

Junos[®] OS Release 12.1X46-D35 for the Branch and High-End SRX Series and J Series

Release 12.1X46-D35
14 July 2015
Revision 2

These release notes accompany Release 12.1X46 of the Junos OS. They describe device documentation and known problems with the software. Junos OS runs on all Juniper Networks SRX Series Services Gateways and J Series Services Routers.

For the latest, most complete information about outstanding and resolved issues with the Junos OS software, see the Juniper Networks online software defect search application at <http://www.juniper.net/prsearch>.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, which is located at <https://www.juniper.net/techpubs/software/junos/>.

Contents

Junos OS Release Notes for Branch SRX Series and J Series	5
New and Changed Features	5
Release 12.1X46-D30 Software Features	6
Release 12.1X46-D20 Software Features	6
Release 12.1X46-D15 Software Features	7
Release 12.1X46-D10 Software Features	8
Changes in Behavior and Syntax	17
Application Firewall	17
Command-Line Interface (CLI)	18
Flow and Processing	21
Hardware	21
Interfaces and Routing	21
Intrusion Detection and Prevention (IDP)	22
J-Web	29
Logical Systems	30
Network Time Protocol	30
Policy Applications	30
Simple Network Management Protocol (SNMP)	30

System Logs	30
System Management	31
User Interface and Configuration	31
Virtual Private Networks (VPNs)	31
Known Behavior	33
Application Layer Gateways (ALGs)	33
AppSecure	34
AX411 Access Points	34
Chassis Cluster	35
Command-Line Interface (CLI)	36
Connectivity Fault Management (CFM)	36
Dynamic Host Configuration Protocol (DHCP)	36
Flow and Processing	37
Hardware	38
Interfaces and Routing	38
Intrusion Detection and Prevention (IDP)	43
IPv6	45
J-Web	45
Layer 2 Transparent Mode	47
Network Address Translation (NAT)	47
Power over Ethernet (PoE)	49
Security Policies	49
Simple Network Management Protocol (SNMP)	50
Switching	50
Unified Access Control	51
Unified Threat Management (UTM)	51
Upgrade and Downgrade	51
USB	51
Virtual Private Networks (VPNs)	51
Known Issues	53
Interfaces and Routing	53
Intrusion Detection and Prevention (IDP)	54
Network Address Translation (NAT)	54
Platform and Infrastructure	54
Resolved Issues	54
Resolved Issues - 12.1X46-D35	55
Resolved Issues - 12.1X46-D30	60
Resolved Issues - 12.1X46-D25	64
Resolved Issues - 12.1X46-D20	66
Resolved Issues - 12.1X46-D15	69
Resolved Issues - 12.1X46-D10	70
Documentation Updates	76
Documentation Updates for the Junos OS Software	
Documentation	77
Documentation Updates for the Junos OS Hardware	
Documentation	83
Migration, Upgrade, and Downgrade Instructions	86
Upgrading and Downgrading among Junos OS Releases	87
Upgrading an AppSecure Device	89

Network and Security Manager Support	89
Upgrade and Downgrade Scripts for Address Book Configuration	89
Hardware Requirements	92
Junos OS Release Notes for High-End SRX Series	95
New and Changed Features	95
Release 12.1X46-D30 Software Features	96
Release 12.1X46-D25 Software Features	96
Release 12.1X46-D20 Software Features	96
Release 12.1X46-D15 Software Features	99
Release 12.1X46-D10 Software Features	99
Hardware Features	111
Changes in Behavior and Syntax	113
Application Firewall	114
Application-Level Distributed Denial of Service	115
Chassis Cluster	115
Command-Line Interface (CLI)	116
Compatibility	119
Flow and Processing	119
Intrusion Detection and Prevention (IDP)	120
J-Web	129
Logical Systems	130
Management Information Bases (MIBs)	130
Network Time Protocol	130
Policy Applications	130
Security Policies	130
Session Timeout for Reroute Failure	131
Simple Network Management Protocol (SNMP)	131
System Logs	132
System Management	133
Unified Threat Management (UTM)	133
Unified In-Service Software Upgrade (ISSU)	134
Virtual Private Networks (VPNs)	134
Known Behavior	135
Application Layer Gateways (ALGs)	135
AppSecure	137
Chassis Cluster	137
Dynamic Host Configuration Protocol (DHCP)	139
Flow and Processing	139
General Packet Radio Service (GPRS)	140
Hardware	143
Interfaces and Routing	143
Intrusion Detection and Prevention (IDP)	145
IP Monitoring	147
IPv6	147
J-Web	148
Logical Systems	149
Network Address Translation (NAT)	150
Security Policies	152
Services Offloading	153

Simple Network Management Protocol (SNMP)	154
Unified Access Control	154
Unified Threat Management (UTM)	154
Virtual Private Networks (VPNs)	154
Known Issues	155
Intrusion Detection and Prevention (IDP)	156
Network Address Translation (NAT)	156
Platform and Infrastructure	156
Unified Threat Management (UTM)	156
Virtual Private Networks (VPN)	156
Resolved Issues	156
Resolved Issues - 12.1X46-D35	157
Resolved Issues - 12.1X46-D30	163
Resolved Issues - 12.1X46-D25	167
Resolved Issues - 12.1X46-D20	171
Resolved Issues - 12.1X46-D15	173
Resolved Issues - 12.1X46-D10	175
Documentation Updates	183
Documentation Updates for the Junos OS Software	
Documentation	183
Documentation Updates for the Junos OS Hardware	
Documentation	190
Migration, Upgrade, and Downgrade Instructions	191
Upgrading and Downgrading among Junos OS Releases	191
Upgrading an AppSecure Device	193
Network and Security Manager Support	194
Upgrade and Downgrade Scripts for Address Book Configuration	194
Upgrade Policy for Junos OS Extended End-Of-Life Releases	196
Hardware Requirements	196
Product Compatibility	197
Hardware Compatibility	197
Finding More Information	197
Documentation Feedback	197
Requesting Technical Support	198
Revision History	200

Junos OS Release Notes for Branch SRX Series and J Series

Powered by Junos OS, Juniper Networks SRX Series Services Gateways provide robust networking and security services. SRX Series Services Gateways range from lower-end branch devices designed to secure small distributed enterprise locations to high-end devices designed to secure enterprise infrastructure, data centers, and server farms. The branch SRX Series Services Gateways include the SRX100, SRX110, SRX210, SRX220, SRX240, SRX550, and SRX650 devices.

Juniper Networks J Series Services Routers running Junos OS provide stable, reliable, and efficient IP routing, WAN and LAN connectivity, and management services for small to medium-sized enterprise networks. These routers also provide network security features, including a stateful firewall with access control policies and screens to protect against attacks and intrusions, and IPsec VPNs. The J Series Services Routers include the J2320, J2350, J4350, and J6350 devices.

- [New and Changed Features on page 5](#)
- [Changes in Behavior and Syntax on page 17](#)
- [Known Behavior on page 33](#)
- [Known Issues on page 53](#)
- [Resolved Issues on page 54](#)
- [Documentation Updates on page 76](#)
- [Migration, Upgrade, and Downgrade Instructions on page 86](#)

New and Changed Features

The following features have been added to Junos OS Release 12.1X46. Following the description is the title of the manual or manuals to consult for further information.



NOTE: For the latest updates about support and issues on Junos Pulse, see the [Junos Pulse Release Notes](#).

- [Release 12.1X46-D30 Software Features on page 6](#)
- [Release 12.1X46-D20 Software Features on page 6](#)
- [Release 12.1X46-D15 Software Features on page 7](#)
- [Release 12.1X46-D10 Software Features on page 8](#)

Release 12.1X46-D30 Software Features

Application Layer Gateways (ALGs)

- **MS-RPC ALG and Sun RPC ALG map table scaling for SRX Series devices**— Starting with Junos OS Release 12.1x46-D30, the MS-RPC ALG and Sun RPC ALG dynamically allocate new mapping entries instead of using a default size (512 entries). They also offer a flexible time-based RPC mapping entry that removes the mapping entry (auto-clean) without affecting the associated active RPC sessions, including both control session and data session.

[See [RPC ALG Feature Guide for Security Devices](#).]

Release 12.1X46-D20 Software Features

Chassis Cluster

- **Autorecovery of fabric link** [SRX Series]—The fabric link feature supports autorecovery, which includes the following enhancements:
 - Fabric monitoring feature is enabled by default on high-end SRX Series, and hence recovery of fabric link and synchronization takes place automatically.
 - If the fabric link goes down, RG1+ becomes ineligible on either the secondary node or the node with failures, by default. The node remains in this state until the fabric link comes up or the other node goes away.
 - If the fabric link goes down followed by the control link, then after approximately 66 seconds the secondary node (or the node with failures) assumes that the remote node is dead and takes over as the primary node.

[See [Understanding Chassis Cluster Fabric Links](#).]

- **Enhanced debugging support for chassis cluster** [SRX Series]—The chassis cluster debugging functionality has the following enhancements:
 - The **show chassis cluster status** command output includes failure reasons (acronyms and their expansions) when the redundancy group's priority is zero.
 - Cleaner jsrpd process includes removing unwanted logs and moving the debug log message from level LOG_INFO to LOG_DEBUG.
 - The **show chassis cluster information** command output displays redundancy group, LED, and monitored failure details.
 - SNMP traps send messages when a node's weight goes down and also when it recovers.
 - The **show chassis cluster ip-monitoring** command output displays both the global threshold and the current threshold of each node and displays the weight of each monitored IP address.
 - A syslog message appears when the control link goes down.

[See [show chassis cluster ip-monitoring status](#).]

Public Key Infrastructure (PKI)

- **Online Certificate Status Protocol (OCSP)** [SRX Series]—OCSP, like CRL, checks the revocation status of X509 certificates. Requests are sent to the OCSP server(s) configured in a CA profile with the **ocsp url** statement at the **[edit security pki ca-profile profile-name revocation-check]** hierarchy level. The **use-ocsp** option must also be configured. If there is no response from the OCSP server, the request is then sent to the location specified in the certificate's AuthorityInfoAccess extension.

[See the “Public Key Infrastructure (PKI)” section in the [Junos OS 12.1X46-D20 Feature Guide](#).]

Routing Protocols

- **OSPFv3 IPsec authentication and confidentiality** [SRX Series]—OSPF for IPv6, also known as OSPF version 3 (OSPFv3), does not have built-in authentication to ensure that routing packets are not altered and re-sent to the router. In Junos OS Release 12.1X46-D20, IPsec can be used to secure OSPFv3 interfaces and virtual links and provide encryption for OSPF packets.

To configure IPsec for OSPF/OSPFv3, define a security association (SA) with the **security-association sa-name** configuration option at the **[edit security ipsec]** hierarchy level. The configured SA is then applied to the OSPF/OSPFv3 interface or virtual link configuration.

[See the “Routing Protocols” section in the [Junos OS 12.1X46-D20 Feature Guide](#).]

Unified Threat Management (UTM)

- **UTM license enforcement** [SRX Series]—License enforcement is supported for UTM features, including Sophos antivirus, enhanced Web filtering, and antispam filtering on all high-end SRX Series devices in addition to branch SRX Series devices. You can add or remove UTM licenses on SRX Series devices. Each feature license is tied to exactly one software feature and is valid for exactly one device.

[Table 1 on page 7](#) lists the license modules and the license names.

Table 1: UTM License Information

UTM Module	License Name
SAV	av_key_sophos_engine
AS	anti_spam_key_sbl
EWf	wf_key_websense_ewf

[See the “UTM” section in the [Junos OS 12.1X46-D20 Feature Guide](#).]

[See [License Enforcement](#).]

Release 12.1X46-D15 Software Features

IP Monitoring

- **IP monitoring with interface as next-hop option [Branch SRX Series]**—IP monitoring enables you to configure a static route with a P2P interface as a next-hop action when IP monitoring has failed.

The following added functions support the track-ip option:

- Next-hop type checking: IP address or interface.
- Interface type checking for next-hop. Only a P2P interface is supported; an error message results when the configuration is committed.
- You can use the interface as a next-hop to construct route parameters and call RPD API to add a static route; log route addition results.
- You can use existing code to delete the route when the primary route recovers.

[See “IP Monitoring” section in [Junos OS 12.1X46-D15 Feature Guide](#).]

Release 12.1X46-D10 Software Features

Application Layer Gateways (ALGs)

- **ALG message buffer optimization**—Starting in Junos OS Release 12.1X46-D10, the ALG message buffer optimization feature has been enhanced to reduce high memory consumption. This feature is supported on all SRX Series and J Series devices.

A message buffer is allocated only when the packet is ready to process. The buffer is freed after the packet completes ALG handling, including modifying the payload, performing NAT, opening a pinhole for a new connection between a client and a server, and transferring data between a client and a server located on opposite sides of a Juniper Networks device.

This feature has the following enhancements:

- Unnecessary objcache buffering is avoided, resulting in low memory utilization.
- jbuf manipulation is used to simplify the message buffer logic.
- Full-fledged message buffer support for the ALG line breaker is more flexible.
- ALG Manager and ALG plug-in logic clarity are optimized.

[See [alg-manager](#).]

- **IPv6 support for SIP ALG**—This feature is supported on all SRX Series and J Series devices.

Starting with Junos OS Release 12.1X46-D10, IPv6 is supported on the SIP ALG along with NAT-PT mode and NAT64 address translation.

The SIP ALG processes the IPv6 address in the same way it processes the IPv4 address for updating the payload if NAT is configured and opening pinholes for future traffic.

NAT-PT is implemented by normal NAT from IPv6 address to IPv4 address and vice versa. The SIP ALG processes those address translations in payload just as the addresses are processed in normal NAT.

NAT64 is a mechanism to allow IPv6 hosts to communicate with IPv4 servers. NAT64 is required to keep the IPv6 to IPv4 address mapping.

Previously, Session Traversal Utilities for NAT (STUN) worked without the SIP ALG. This means that the SIP ALG was not involved when persistent NAT was configured.

Starting with Junos OS Release 12.1X46-D10, STUN can coexist with the SIP ALG and SIP ALG is involved when persistent NAT is configured.

[See [SIP ALG Feature Guide for Security Devices](#).]

- **IPv6 support for RTSP ALG**—This feature is supported on all SRX Series and J Series devices.

Real-Time Streaming Protocol (RTSP) is an Application Layer protocol for controlling the delivery of data with real-time properties. The RTSP ALG accesses existing media files over the network and controls the replay of the media.

Starting with Junos OS Release 12.1X46-D10, IPv6 is supported on the RTSP ALG along with NAT-PT mode and NAT64 address translation.

This feature enables the RTSP ALG to parse IPv6 RTSP packets, open an IPv6 pattern pinhole, and translate the Layer 7 IPv6 address according to the NAT configuration. Also, support for IPv6 RTSP transaction pass through under permission policy and IPv6 RTSP transaction pass through under NAT-PT and NAT 64 are enabled.

[See [SIP RTSP ALG Feature Guide for Security Devices](#).]

- **IPv6 support for PPTP ALG**—Starting with Junos OS Release 12.X46-D10, this feature is supported on all SRX Series devices.

PPTP ALG provides an ALG for the Point-to-Point Tunneling Protocol (PPTP). The PPTP is a Layer 2 protocol that tunnels PPP data across TCP/IP networks. The PPTP client is freely available on Windows systems and popularly applied on Linux systems; it is widely deployed for building VPNs.

To support IPv6, the PPTP ALG parses both IPv4 and IPv6 PPTP packets, performs NAT, and then opens a pinhole for the data tunnel. The flow module supports IPv6 to parse the GRE packet and use the GRE call ID as fake port information to search the session table and gate table.

- **Support for SCCP v20**—This feature is supported on all SRX Series devices.

Starting in Junos OS Release 12.1X46-D10, the SCCP ALG supports SCCP versions 16, 17, and 20 and several SCCP messages have been updated with a new format. Cisco Call Manager (CM) version 7 uses SCCP version 20.

[See [SCCP ALG Feature Guide for Security Devices](#).]

AppSecure

- **Application-aware quality of service (AppQoS)**—Starting in Junos OS Release 12.1X46-D10, AppQoS is supported on all branch SRX Series devices.

AppQoS provides a mechanism for prioritizing traffic utilizing the results of the Application Identification Engine. AppQoS provides application-level traffic control for administrators needing to ensure that business-critical applications get preferential treatment.

AppQoS enables the network administrator to meter, mark, and honor traffic priority based on application policies. It provides application-aware DSCP marking by implementing Layer 7 application-based DSCP rewriters. To apply different loss priority levels to different traffic groups, Layer 2-based to Layer 4-based honoring has been expanded to Layer 7. AppQoS accomplishes application-aware rate limiting by setting the bandwidth limit and burst size limit for different applications.

[See [Understanding Application QoS \(AppQoS\)](#).]

Dynamic Host Configuration Protocol (DHCP)

- **DHCP relay**—Starting in Junos OS Release 12.1X46-D10, the existing DHCP relay feature on all branch SRX Series devices has been enhanced to include chassis cluster support.

[See [Understanding DHCP Relay Agent Operation](#).]

Flow and Processing

- **Enhanced IPv6 support for the screen feature**—This feature is supported on all branch SRX Series and J Series devices.

IPv6 support is extended for the following screen features:

- IPv6 extension header checking and filtering
- IPv6 packet header checking and filtering
- ICMPv6 checking and filtering

New statements and commands allow you to configure these enhancements using security zones similar to previous screen configurations. You can enable, disable, and update screens to drop packets, create logs, and provide increased statistics for IPv6 traffic.



NOTE: By default, IPv6 packets bypass the screen feature.

[See [Understanding IPv6 Support for Screens](#).]

- **Enhanced IPv6 support for flow**—This feature is supported on all branch SRX Series and J Series devices.

IPv6 support is extended for checking and filtering IPv6 extension headers (in accordance with RFC 2460) and IPv6 link-local addresses (in accordance with RFC 4291) in a flow. Nonconforming IPv6 packets will be discarded.

- **Enhancements to flow trace options**—This feature is supported on all branch SRX Series and J Series devices.

Starting in Junos OS Release 12.1X46-D10, flow trace granularity has been enhanced to filter logs effectively. As a result you can access relevant trace messages easily and avoid large traces that slow down your system. You can set the level of message you want displayed by using the new **trace-level** statement at the **[edit security flow traceoptions]** hierarchy level. You can use new flags to trace additional operations such as fragmentation, high availability, multicast, session, tunnel, and route.

[See [traceoptions \(Security Flow\)](#).]

- **Monitoring flow sessions**—This feature is supported on all branch SRX Series and J Series devices.

Beginning with Junos OS Release 12.1X46-D10, you can monitor flow using filters that match different criteria (such as source and destination addresses). New operational mode commands **monitor security flow filter** and **monitor security flow file** have been added. These commands allow you to debug without having to commit or modify your running configuration. Previously, you were required to commit the configuration to turn on trace options, which could possibly change the state of your device.

[See [Monitoring Security Flow Sessions Overview](#).]

Intrusion Detection and Prevention (IDP)

- **IDP IPv6 inspection**—Starting in Junos OS Release 12.1X46-D10, IDP supports IPv6 inspection on the SRX100, SRX210, SRX220, SRX240, SRX550, and SRX650. IPv6 builds upon the functionality of IPv4, providing improvements to addressing, configuration and maintenance, and security.

This feature supports:

- IPv6 traffic inspection
- Attack detection inspection in protocol decoders that support IPv6
- IDP signature database
- IDP logging
- Application identification results

Use the **show security flow session idp family** command with the **inet** or **inet6** option to view IPv4 or IPv6 statistics.

[See [IDP Monitoring and Troubleshooting Guide for Security Devices](#).]

- **IDP security packet capture**—Starting in Junos OS Release 12.1X46-D10, this feature is supported on the SRX100, SRX210, SRX220, SRX240, SRX550, and SRX650.

Viewing packets that precede and follow an attack helps you determine the purpose and extent of an attempted attack, whether an attack was successful, and if any network damage was caused. Packet analysis also aids in defining attack signatures to minimize false positives.

Use the **show security idp counters packet-log** command to display details about the progress, success, and failure of packet capture activity.

You can specify pre-attack, post-attack, and post-attack timeout values. The pre-attack and post-attack default values are 1, and the default post-attack timeout value is 5.



NOTE: Support for packet capture is available only once on each session.

[See [Understanding Security Packet Capture](#).]

IP Spoofing

- **IP spoofing in transparent mode**—Starting in Junos OS Release 12.1X46-D10, this feature is supported on all branch SRX Series devices.

The IP spoofing feature has been enhanced to include Layer 2 transparent mode support. IP spoofing is most frequently used in denial-of-service attacks. In an IP spoofing attack, the attacker gains access to a restricted area of the network and inserts a false source address in the packet header to make the packet appear to come from a trusted source. When SRX Series devices are operating in transparent mode, the IP spoof-checking mechanism makes use of address book entries.



NOTE:

- IP spoofing in Layer 2 transparent mode does not support DNS and wildcard addresses.
- IP spoofing in Layer 2 transparent mode is not supported on IPv6, because branch SRX Series devices do not support IPv6 in Layer 2 transparent mode.

[See [Understanding IP Spoofing in Layer 2 Transparent Mode](#).]

J-Web

- **Management support for NAT options**—Starting in Junos OS Release 12.1X46-D10, support is provided to monitor the following NAT options on all SRX Series devices:
 - Utilization for all source pools
 - Successful, failed, and current sessions for source pools, source rules, destination rules, and static rules
 - Source addresses and source ports for static rules
 - Source ports for source rules
- Support is provided to configure the following NAT options on all SRX Series devices:
 - Source address and port as match criteria for static rules
 - Source port as match criteria for source rules
 - Upper and lower thresholds at which an SNMP trap is triggered for source rules and pools, destination rules, and static rules

- **User firewall J-Web support**

- **Source identity-based firewall policy**—Starting in Junos OS Release 12.1X46-D10, this feature is supported on the existing Firewall Policies Configuration and Monitoring Policies pages on all branch SRX Series devices. This feature allows you to configure and monitor source identities in a firewall policy.
- **New J-Web pages for user firewall**—Starting in Junos OS Release 12.1X46-D10, new user firewall pages are supported on all branch SRX Series devices.

The following webpages have been added to the J-Web user interface:

- **Authentication Priority Configuration Page**—You can either disable an optional authentication source or reassign a unique priority to it.
 - **Local Authentication Configuration Page and Local Authentication Monitoring Page**—You can configure and monitor local Firewall authentication.
 - **UAC Settings Configuration Page and UAC Authentication Monitoring Page**—You can configure UAC and monitor UAC authentication.
- **Allow adding a new policy and moving an existing policy to an arbitrary location**
 - **Firewall Policies Configuration Page Options**—Starting in Junos OS Release 12.1X46-D10, several new options on the Firewall Policies Configuration page are supported on all branch SRX Series devices. The Add menu includes Add before and Add after options that allow you to add a new policy before or after a selected policy. On the Move menu, there is a new Move to option that allows you to specify a target location. You can also drag and drop a policy to the target location.
 - **Checking Policies Monitoring Page**—Starting in Junos OS Release 12.1X46-D10, the Move to option on the Checking Policies Monitoring page is supported on all branch SRX Series devices.

Management Information Bases (MIBs)

- **SNMP aggregation for policy MIBs**—Starting in Junos OS Release 12.1X46-D10, this feature is supported on all SRX Series devices.

A set of systemwide policy statistics such as policy-allowed packets, bytes and rates, policy-dropped packets, bytes and rates, policy flows allowed, and rate statistics have been added in the enterprise-specific policy MIB JUNIPER-JS-POLICY-MIB. You can obtain the policy statistics by using the SNMP agent or the CLI operational mode commands. Use the following CLI commands to set, clear, and display the systemwide policy statistics:

- **set security policies policy-stats system-wide <disable | enable>**—Configures systemwide policy statistics. Disabled by default.
- **clear security policies statistics**—Clears the systemwide policy statistics.
- **show snmp mib walk jnxJsPolicySystemStats**—Displays both IPv4 and IPv6 statistics.
- **show snmp mib walk jnxJsPolicySystemStatsIPv4**—Displays only IPv4 statistics.

[See [Policy Objects MIB](#).]

Virtual Private Networks (VPNs)

- **Enhanced X2 interface monitoring**—This feature is supported on all SRX Series devices.

In an LTE mobile network, X2 interfaces are used to connect Evolved Node Bs (eNodeBs) for signal handover, monitoring, and radio coverage. SRX Series devices connect these eNodeBs using IPsec tunnels.

This feature enables you to monitor traffic between eNodeBs by snooping into the clear text traffic as it flows from one IPsec tunnel to another. Use the **monitor-filter** statement at the **[edit security forwarding-options]** hierarchy level to duplicate clear text packets and send them to the physical interface. You can then use Ethereal or other packet analyzers to verify or collect the X2 traffic.

[See [Understanding X2 Traffic Monitoring](#)]

- **Support for IPv6 address encapsulation in route-based one-to-one site-to-site VPN tunnels**—This feature is supported on all SRX Series devices.

In tunnel mode, IPsec encapsulates the original IP datagram—including the original IP header—within a second IP datagram. The outer IP header contains the IP address of the gateway, while the inner header contains the ultimate source and destination IP addresses. The outer and inner IP headers can have a protocol field of IPv4 or IPv6. As of Junos OS Release 12.1X46-D10, the following tunnel modes are supported on SRX Series devices:

- IPv4-in-IPv4 tunnels encapsulate IPv4 packets inside IPv4 packets.
- IPv6-in-IPv6 tunnels encapsulate IPv6 packets inside IPv6 packets.
- IPv6-in-IPv4 tunnels encapsulate IPv6 packets inside IPv4 packets.
- IPv4-in-IPv6 tunnels encapsulate IPv4 packets inside IPv6 packets.

There are no new CLI configuration statements for this feature.

IPv4 and IPv6 traffic can be routed into a single IPv4 or IPv6 tunnel; the st0 interface bound to the tunnel must be configured for both family inet and family inet6. Dual stack tunnels—parallel IPv4 and IPv6 tunnels over a single physical external interface to different VPN peers—are also supported.

[See [VPN Feature Support for IPv6 Addresses](#).]

- **Dead peer detection (DPD) enhancements**—This feature is supported on all SRX Series devices.

Network devices use the DPD protocol to verify the existence and availability of other peer devices. The default DPD mode **optimized** sends probes if there is no incoming IKE or IPsec traffic from the peer within a configured interval after outgoing packets are sent to the peer. The **always-send** option sends DPD probes at configured intervals regardless of traffic activity between peers. A new configuration option **probe-idle-tunnel** at the **[edit security ike gateway dead-peer-detection]** hierarchy level sends DPD probes when there is no incoming or outgoing IKE or IPsec traffic between peers.



NOTE: We recommend that you configure **probe-idle-tunnel** instead of **always-send**.

For all DPD modes, Phase 1 and Phase 2 security associations are cleared if a specified number of probes are sent with no response from the peer.

[See [Understanding Dead Peer Detection](#).]

- **Multiple traffic selectors on a route-based VPN**—This feature is supported on all branch SRX Series devices.

A traffic selector (also known as a proxy ID in IKEv1) is an agreement between IKE peers to permit traffic through a tunnel if the traffic matches a specified pair of local and remote addresses. With this feature, you can define multiple traffic selectors within a specific route-based VPN, resulting in a unique SA for each traffic selector configured. Only traffic that conforms to a traffic selector is permitted through the associated IPsec SA.

To configure a traffic selector, use the **traffic-selector** configuration statement at the **[edit security ipsec vpn vpn-name]** hierarchy level. The traffic selector pair is defined with the mandatory **local-ip ip-address** and **remote-ip ip-address** statements. The CLI operational command **show security ipsec security-association traffic-selector traffic-selector** displays SA information for the specified traffic selector.

[See [Understanding Traffic Selectors in Route-Based VPNs](#).]

- **IKEv2 configuration payload support with RADIUS**—This feature is supported on all SRX Series devices.

Configuration payload is an Internet Key Exchange (IKE) version 2 feature used to propagate provisioning information from an IKE responder to the IKE initiator. Starting with Junos OS Release 12.1X46-D10, IKEv2 configuration payload is supported with route-based VPNs only. The following attribute types, defined in RFC 5996, *Internet Key Exchange Protocol Version 2 (IKEv2)*, can be returned to the IKE initiator by the IKE responder:

- INTERNAL_IP4_ADDRESS
- INTERNAL_IP4_NETMASK
- INTERNAL_IP4_DNS

For the IKE responder to provide the initiator with provisioning information, it must acquire the information from a specified source such as a RADIUS server. Provisioning information can also be returned from a DHCP server through a RADIUS server. On the RADIUS server, the user information should not include an authentication password. As in previous Junos OS releases for the SRX Series, the RADIUS server profile is bound to the IKE gateway using the **xauth access-profile profile-name** configuration at the **[edit security ike gateway gateway-name]** hierarchy level.

This feature is supported only for point-to-multipoint secure tunnel (st0) interfaces. For point-to-multipoint interfaces, the interfaces must be numbered and the addresses

in the configuration payload INTERNAL_IP4_ADDRESS attribute type must be within the subnet range of the associated point-to-multipoint interface.



NOTE: IKEv2 on SRX Series devices does not support policy-based VPNs or VPN monitoring.

[See [Understanding Internet Key Exchange Version 2.](#)]

- **IKEv2 with NAT-T and dynamic endpoint VPN**—This feature is supported on all SRX Series devices.

Starting with Junos OS 12.1X46-D10, both IKEv2 initiators and responders in a route-based VPN can be behind NAT devices. The IKEv2 NAT-T feature supports IPsec traffic that crosses NAT devices. Static NAT and dynamic NAT are supported. In static NAT, there is a one-to-one relationship between the private and the public addresses. In dynamic NAT, there is a many-to-one or many-to-many relationship between the private and public addresses.

Dynamic endpoint (DEP) VPN is a Junos OS feature that covers IKEv2 initiator and responder perspectives. From the initiator's perspective, DEP VPN covers the situation where the IKE external interface address is not fixed and is therefore not known by the responder. This situation can occur when the peer's address is dynamically assigned by an ISP or when the peer's connection crosses a NAT device that allocates addresses from a dynamic address pool. From the responder's perspective, DEP VPN describes either a finite number of VPNs that are created for a number of VPN peers in a many-to-many scenario or a shared VPN in a many-to-one scenario.

Starting with Junos OS 12.1X46-D10, the default value for the **nat-keepalive** option configured at the **[edit security ike gateway gateway-name]** hierarchy level has been changed from 5 seconds to 20 seconds.

[See [Understanding NAT-T.](#)]

Web Authentication

- **Web-redirect firewall authentication**—Starting in Junos OS Release 12.1X46-D10, Web authentication redirect enhancement is provided on all SRX Series devices.

With this feature, when you attempt to initiate a connection across the firewall, after successful authentication the browser launches your original destination URL without you needing to retype the URL.

The following message is displayed:

Redirecting to the original url, please wait

[See [Firewall User Authentication Overview](#)]

Related Documentation

- [Changes in Behavior and Syntax on page 17](#)
- [Known Behavior on page 33](#)
- [Known Issues on page 53](#)

- [Resolved Issues on page 54](#)
- [Documentation Updates on page 76](#)
- [Migration, Upgrade, and Downgrade Instructions on page 86](#)

Changes in Behavior and Syntax

The following current system behavior, configuration statement usage, and operational mode command usage might not yet be documented in the Junos OS documentation:

Application Firewall

- Prior to Junos OS Release 12.1X46-D10, when a rule specifies **dynamic-application junos:HTTP** without specifying any other nested application, the rule matches all HTTP traffic whether the traffic contains a nested application or not.

In Junos OS Release 12.1X46-D15 and later, that functionality has changed. When a rule specifies **dynamic-application junos:HTTP**, only HTTP traffic with no nested members is matched.

Consider the following application firewall ruleset:

```
rule-sets http-ruleset {
  rule rule1 {
    match {
      dynamic-application [junos:HTTP];
    }
    then {
      deny;
    }
  }
  default-rule {
    permit;
  }
}
```

Prior to Junos OS Release 11.4R6, the sample rules would be applied to traffic as shown in the following list:

- HTTP traffic with or without nested applications would be denied by rule1.
HTTP traffic with a nested application, such as junos:FACEBOOK or junos:TWITTER, would be denied by rule1.
- All other traffic would be permitted by the default rule.

In Junos OS Release 11.4R6 and later, the dynamic application junos:HTTP matches only the HTTP traffic that contains no recognizable nested application. The sample rules would now be applied differently:

- Only the HTTP traffic with no nested application would be denied by rule1.
HTTP traffic with a nested application, such as junos:FACEBOOK or junos:TWITTER, would no longer match rule1.
- All other traffic would be permitted by the default rule.

HTTP traffic with a nested application, such as `junos:FACEBOOK` or `junos:TWITTER`, would be permitted by the default rule.

- In Junos OS Release 12.1X46-D10 and earlier, if a nested application is not configured in any rule, then the nested application would match the default rule and take action specified in the default rule.

Starting in Junos OS Release 12.1X46-D10, the functionality has changed. If a nested application matches the default rule, then the application firewall uses the application type to match the rule and takes action specified in the rule. Use the **`set security application-firewall nested-application dynamic-lookup enable`** command to control the behavior of the nested application, so that both the application and the nested application are consistent.

The default behavior of nested application before Junos OS Release 12.1X46-D10:

- Application firewall matches with the specific rule, if the nested application is configured explicitly in a rule.
- Application firewall matches with the default rule, if the nested application is not configured explicitly in a rule.
- Records the statistics of the application firewall in the matched rule.

The new behavior of nested application in Junos OS Release 12.1X46-D10:

- Application firewall matches with an application rule during application firewall policy lookup, if there is no explicit rule for the nested application.
- Application firewall matches with a specific rule, if the nested application is configured explicitly in a rule.
- Records the statistics of the application firewall in the matched rule.

Command-Line Interface (CLI)

New or Changed CLI

- Starting in Junos OS Release 12.1X46-D30, for all branch SRX Series devices there is an option to remove the peer loop check for private AS numbers. The **`no-peer-loop-check`** option has been added under the **`remove-private`** command at the following hierarchy levels:

```
[edit logical-systems logical-system-name protocols bgp]
```

```
[edit protocols bgp]
```

```
[edit routing-instances routing-instance-name protocols bgp]
```

- Starting in Junos OS Release 12.1X46-D20, for all branch SRX Series devices in chassis cluster mode, there is a **`node`** option available for all **`show chassis`** CLI commands. The **`node`** option displays status information for all FPCs or for the specified FPC on a specific node (device) in the cluster.
- Prior to Junos OS Release 12.1X46-D10, when you configured the DNS proxy server using the **`set system services dns dns-proxy view view-name domain domain-name forwarder`** CLI statement, if the IP address specified in the forwarder option was not

available, the DNS query was forwarded to the default DNS servers (DNS servers provided by the ISP). The device acquired the public IP addresses from the default DNS servers.

Starting in Junos OS Release 12.1X46-D10, the **forward-only** option is added to the **set system services dns dns-proxy view view-name domain domain-name forward-only** CLI statement.

You can use the **forward-only** option to prevent the device from acquiring the public IP addresses from the DNS servers (by terminating the DNS query) in cases when the specified IP address is unreachable.

- On all branch SRX Series and J Series devices, the following commands are now supported:

CLI Command	Description
show pppoe interfaces	List all PPPoE sessions.
request pppoe connect	Connect to all sessions that are down.
request pppoe connect <i>pppoe interface name</i>	Connect only to the specified session.
request pppoe disconnect	Disconnect all sessions that are up.
request pppoe disconnect <i>session id or pppoe interface name</i>	Disconnect only the specified session, identified by either a session ID or a PPPoE interface name.

- On all J Series devices, a new CLI **request system (halt | power-off | reboot) power-off fpc** command has been introduced to bring Flexible PIC Concentrators (FPCs) offline before Routing Engines are shut down. This command prevents the short network outage because of the Layer 2 loop.

CLI Command	Description
request system halt power-off fpc	Bring FPC offline and then halt the system.
request system power-off power-off fpc	Bring FPC offline and then power off the system.
request system reboot power-off fpc	Bring FPC offline and then reboot the system.

Deprecated Items for Security Hierarchy

- [Table 2 on page 20](#) lists deprecated items (such as CLI statements, commands, options, and interfaces).

CLI statements and commands are deprecated—rather than immediately removed—to provide backward compatibility and a chance to bring your configuration into compliance with the new configuration. We strongly recommend that you phase out deprecated items and replace them with supported alternatives.

Table 2: Items Deprecated in Release 12.1

Deprecated Item	Replacement	Hierarchy Level or Command Syntax	Additional Information
<code>download-timeout</code>	-	<code>download-timeout timeout</code>	On all branch SRX Series devices, the download-timeout command is deprecated. If the configuration is present, then that configuration will be ignored. The IDP process internally triggers the security package to install when an automatic download is completed. There is no need to configure any download timeout.
<code>node</code>	-	<code>request security idp security-package download</code>	On all branch SRX Series devices operating in a chassis cluster, the request security idp security-package download command with the node option is not supported: <code>request security idp security-package download node primary</code> <code>request security idp security-package download node local</code> <code>request security idp security-package download node all</code>

Compatibility

- **Version compatibility for Junos SDK**—Beginning with Junos OS Release 12.1X44-D10, Junos OS applications will install on the Junos OS only if the application is built with the same release as the Junos OS release on which the application is being installed.

For example, an application built with Junos OS Release 12.1R2 will only install on Junos OS Release 12.1R2 and will not install on Junos OS Release 12.1R1 or Junos OS Release 12.1R3.

Flow and Processing

- The minimum value you can configure for TCP session initialization is 4 seconds. The default value is 20 seconds; if required you can set the TCP session initialization value to less than 20 seconds.
- On all branch SRX Series devices, the default value of type of service (ToS) for IKE packets has been changed from 0x00 to 0xc0.
- On all branch SRX Series and J Series devices, you can configure the TCP session timeout in a half-closed state by using the **apply-to-half-close-state** statement at the **[edit security flow tcp-session time-wait-state]** hierarchy level. This enables the system to apply the configured session timeout on receiving only one FIN packet (either client-to-server or server-to-client). When this statement is not configured, the default behavior takes effect, which is to apply the configured session timeout on receiving both the FIN packets. The default TCP session timeout remains 150 seconds. [See *apply-to-half-close-state*.]

Hardware

- On SRX550 devices, the mini-USB console cable provides a “break” message to the Windows application whenever the console cable is unplugged and re-plugged. If you have configured “debugger-on-break”, the system goes to the **db>** prompt because the system receives a break character. This behavior is specific to the mini-USB console.
- Starting in Junos OS Release 12.1X46-D15, external clocking is enabled on SRX550 devices with a DS3/E3 interface. In Junos OS Release 12.1X46-D10 and earlier, the external clocking option was disabled to overcome the limitations present in the hardware to support this clocking option.

Interfaces and Routing

- A new attribute, **max-synacks-queued**, is added to IDP sensor configuration TCP reassembler. This attribute defines the maximum syn/ack queued with different SEQ numbers and takes the values 0 through 5. Also, a new counter, **Duplicate Syn/Ack with different SEQ**, is added to the IDP TCP reassembler. This counter displays the number of syn/ack packets with different SEQ numbers.
- On SRX240 and SRX650 devices, for the Layer 2 LAG interface, the hash algorithm for load balancing is now based on source IP address and destination IP address instead of source MAC address and destination MAC address.

Intrusion Detection and Prevention (IDP)

- In Junos OS releases earlier than Junos OS Release 12.1X46-D25, TACACS+ options for authentication and accounting did not include an option for configuring a timestamp and time zone.

In Junos OS Release 12.1X46-D25 and later releases, you can use the **timestamp-and-timezone** option at the **[edit system tacplus-options]** hierarchy to include start time, stop time, and time zone attributes in start/stop accounting records. [See *tacplus-options*.]

- A system log message is generated when an IDP signature database update or policy compilation fails with an empty dynamic group. The system-generated log message is **Dynamic Attack group [dyn_group_1] has no matching members found. Group is empty.**
- By default, values for IDP reassembler packet memory and application identification packet memory used by IDP are established as percentages of all memory. In most cases, these default values are adequate.
- If a deployment exhibits an excessive number of dropped TCP packets or retransmissions resulting in high IDP reassembly memory usage, use the following option:

The **max-packet-mem-ratio** option to reset the percentage of available IDP memory for IDP reassembly packet memory. Acceptable values are between 5 and 40 percent.

```
set security idp sensor-configuration re-assembler max-packet-mem-ratio  
percentage-value
```



NOTE: The **max-packet-mem** option has been deprecated and replaced by the new **max-packet-mem-ratio** option.

- If a deployment exhibits an excessive number of ignored IDP sessions due to reassembler and application identification memory allocation failures, use the following options:
 - The **max-packet-memory-ratio** option sets application identification packet memory limit as a percentage of available IDP memory. This memory is only used by IDP in cases where application identification delays identifying an application. Acceptable values are between 5 and 40 percent.

```
set security idp sensor-configuration application-identification  
max-packet-memory-ratio percentage-value
```
 - The **max-reass-packet-memory-ratio** option sets the reassembly packet memory limit for application identification as a percentage of available IDP memory. Acceptable values are between 5 and 40 percent.

```
set security idp sensor-configuration application-identification  
max-reass-packet-memory-ratio percentage-value
```



NOTE: The `max-packet-memory` option has been deprecated and replaced by the new `max-packet-memory-ratio` and `max-reass-packet-memory-ratio` options.

- On all branch SRX Series devices with a single session, when IDP is activated, the upload and download speeds are slow when compared to the firewall performance numbers.

To overcome this issue, a new CLI command, `set security idp sensor-configuration ips session-pkt-depth`, is introduced, for which the `session-pkt-depth sensor-configuration` value is global for any session.

The `session-pkt-depth sensor-configuration` value specifies the number of packets per session that are inspected by IDP. Any packets beyond the specified value are not inspected. For example, when `session-pkt-depth sensor-configuration` is configured as “n”, the IDP inspection happens only for first (n-1) packets in that session. Packets from the nth packet onwards are ignored by IDP.

The default value of `session-pkt-depth sensor-configuration` is zero. When the default value of zero is used, the session-pkt-depth value is not addressed, and IDP performs a full inspection of the session.

- Starting in Junos OS Release 12.1X46-D25, the `show security idp counters flow` command output is changed to include new fields.

[Table 3 on page 23](#) lists the output fields for the `show security idp counters flow` command. Output fields are listed in the approximate order in which they appear.

Table 3: show security idp counters flow Output Fields

Field Name	Description
Fast-path packets	Number of packets that are set through fast path after completing IDP policy lookup.
Slow-path packets	Number of packets that are sent through slow path during IDP policy lookup.
Session construction failed (Unsupported)	Number of times the packet failed to establish the session.
Session limit reached	Number of sessions that reached IDP sessions limit.
Session inspection depth reached	Number of sessions that reached inspection depth.
Memory limit reached	Number of sessions that reached memory limit.
Not a new session (Unsupported)	Number of sessions that extended beyond time limit.

Table 3: show security idp counters flow Output Fields (*continued*)

Field Name	Description
Invalid index at age-out (Unsupported)	Invalid session index in session age-out message.
Packet logging	Number of packets saved for packet logging.
Policy cache hits	Number of sessions that matched policy cache.
Policy cache misses	Number of sessions that did not match policy cache.
Policy cache entries	Number of policy cache entries.
Maximum flow hash collisions	Maximum number of packets, of one flow, that share the same hash value.
Flow hash collisions	Number of packets that share the same hash value.
Gates added	Number of gate entries added for dynamic port identification.
Gate matches (Unsupported)	Number of times a gate is matched.
Sessions deleted	Number of sessions deleted.
Sessions aged-out (Unsupported)	Number of sessions that are aged out if no traffic is received within session timeout value.
Sessions in-use while aged-out (Unsupported)	Number of sessions in use during session age-out.
TCP flows marked dead on RST/FIN	Number of sessions marked dead on TCP RST/FIN.
policy init failed	Policy initiation failed.
Number of sessions exceeds high mark	Number of sessions that exceed high mark.
Number of sessions drops below low mark	Number of sessions that fall below low mark.
Memory of sessions exceeds high mark	Session memory exceeds high mark.
Memory of sessions drops below low mark	Session memory drops below low mark.

Table 3: show security idp counters flow Output Fields (*continued*)

Field Name	Description
Sessions constructed	Number of sessions established.
SM Sessions encountered memory failures	Number of SM sessions encountered memory failure.
SM Packets on sessions with memory failures	Number of SM packets on SM sessions with memory failure.
SM Sessions dropped	Number of SM sessions dropped.
SM sessions ignored	Number of sessions ignored in Security Module (SM).
SM sessions interested	Number of SM sessions interested.
SM sessions not interested	Number of SM sessions not interested.
SM sessions interest error	Number of errors created for SM sessions interested.
Sessions destructed	Number of sessions destructed.
SM Session Create	Number of SM sessions created.
SM Packet Process	Number of packets processed from SM.
SM FTP data session ignored by IDP	Number of SM FTP data sessions that are ignored by IDP.
SM Session close	Number of SM sessions closed.
SM client-to-server packets	Number of SM client-to-server packets.
SM server-to-client packets	Number of SM server-to-client packets.
SM client-to-server L7 bytes	Number of SM client-to-server Layer 7 bytes.
SM server-to-client L7 bytes	Number of SM server-to-client Layer 7 bytes.
Client-to-server flows ignored	Number of client-to-server flow sessions that are ignored.
Server-to-client flows ignored	Number of server-to-client flow sessions that are ignored.
Both directions flows ignored	Number of server-to-client and client-to-server flow sessions that are ignored.
Fail-over sessions dropped	Number of fail-over sessions dropped.

Table 3: show security idp counters flow Output Fields (*continued*)

Field Name	Description
Sessions dropped due to no policy	Number of sessions dropped because there was no active IDP policy.
IDP Stream Sessions dropped due to memory failure	Number of IDP stream sessions that are dropped because of memory failure.
IDP Stream Sessions ignored due to memory failure	Number of IDP stream sessions that are ignored because of memory failure.
IDP Stream Sessions closed due to memory failure	Number of IDP stream sessions that are closed because of memory failure.
IDP Stream Sessions accepted	Number of IDP stream sessions that are accepted.
IDP Stream Sessions constructed	Number of IDP stream sessions that are constructed.
IDP Stream Sessions destructed	Number of IDP stream sessions that are destructed.
IDP Stream Move Data	Number of Stream data events handled by IDP.
IDP Stream Sessions ignored on JSF SSL Event	Number of IDP stream sessions that are ignored because of a JSF SSL proxy event.
IDP Stream Sessions not processed for no matching rules	Number of IDP stream sessions that are not processed for no matching rules.
IDP Stream stbuf dropped	Number of IDP stream plugin buffers dropped.
IDP Stream stbuf reinjected	Number of IDP stream plugin buffers injected.
Busy packets from stream plugin	Number of packets saved as one or more packets of this session from stream plugin.
Busy packets from packets plugin	Number of saved packets for IDP stream plugin sessions.
Bad kpp	Number of internal marked packets logged for IDP processing.
Lsys policy id lookup failed sessions	Number of sessions that failed logical systems policy lookup
Busy packets	Number of packets saved as one or more packets of this session are handed off for asynchronous processing.

Table 3: show security idp counters flow Output Fields (*continued*)

Field Name	Description
Busy packet errors	Number of packets found with IP checksum error after asynchronous processing is completed.
Dropped queued packets (async mode)	Number of queued packets dropped based on policy action, reinjection failures, or if the session is marked to destruct.
Dropped queued packets failed (async mode)	Not used currently.
Reinjected packets (async mode)	Number of packets reinjected into the queue.
Reinjected packets failed (async mode)	Number of failed reinjected packets.
AI saved processed packet	Number of AI packets saved for which the asynchronous processing is completed.
Busy packet count incremented	Number of times the busy packet count incremented in asynchronous processing.
busy packet count decremented	Number of times the busy packet count decremented in asynchronous processing.
session destructed in pme	Number of sessions destructed as a part of asynchronous result processing.
session destruct set in pme	Number of sessions set to be destructed as a result of asynchronous processing.
KQ op	Number of sessions with one of the following status: <ul style="list-style-type: none"> • KQ op hold—number of times packets held by IDP. • KQ op drop—number of times packets dropped by IDP. • KQ op route—number of times IDP decided to be route the packet directly. • KQ op Continue—number of times IDP decided to continue to process the packet. • KQ op error—number of times error occurred while IPD processing packet. • KQ op stop—number of times IDP decided to stop processing the packet.
PME wait not set	Number of AI saved packets given for signature matching.
PME wait set	Number of packets given for signature matching without AI save.

Table 3: show security idp counters flow Output Fields (*continued*)

Field Name	Description
PME KQ run not called	Number of times signature matching results processed out of packet receiving order.

```
user@host> show security idp counters flow
```

IDP counter type	Value
Fast-path packets	0
Slow-path packets	0
Session construction failed	0
Session limit reached	0
Session inspection depth reached	0
Memory limit reached	0
Not a new session	0
Invalid index at ageout	0
Packet logging	0
Policy cache hits	0
Policy cache misses	0
Maximum flow hash collisions	0
Flow hash collisions	0
Gates added	0
Gate matches	0
Sessions deleted	0
Sessions aged-out	0
Sessions in-use while aged-out	0
TCP flows marked dead on RST/FIN	0
Policy init failed	0
Number of times Sessions exceed high mark	0
Number of times Sessions drop below low mark	0
Memory of Sessions exceeds high mark	0
Memory of Sessions drops below low mark	0
SM Sessions encountered memory failures	0
SM Packets on sessions with memory failures	0
Sessions constructed	0
SM Sessions ignored	0
SM Sessions dropped	0
SM Sessions interested	0
SM Sessions not interested	0
SM Sessions interest error	0
Sessions destructed	0
SM Session Create	0
SM Packet Process	0
SM ftp data session ignored by idp	0
SM Session close	0
SM Client-to-server packets	0
SM Server-to-client packets	0
SM Client-to-server L7 bytes	0
SM Server-to-client L7 bytes	0
Client-to-server flows ignored	0
Server-to-client flows ignored	0
Both directions flows ignored	0
Fail-over sessions dropped	0
Sessions dropped due to no policy	0
IDP Stream Sessions dropped due to memory failure	0
IDP Stream Sessions ignored due to memory failure	0
IDP Stream Sessions closed due to memory failure	0
IDP Stream Sessions accepted	0

IDP Stream Sessions constructed	0
IDP Stream Sessions destructed	0
IDP Stream Move Data	0
IDP Stream Sessions ignored on JSF SSL Event	0
IDP Stream Sessions not processed for no matching rules	0
IDP Stream stbuf dropped	0
IDP Stream stbuf reinjected	0
Busy pkts from stream plugin	0
Busy pkts from pkt plugin	0
bad kpp	0
Lsys policy id lookup failed sessions	0
Busy packets	0
Busy packet Errors	0
Dropped queued packets (async mode)	0
Dropped queued packets failed(async mode)	0
Reinjected packets (async mode)	0
Reinjected packets failed(async mode)	0
AI saved processed packet	0
busy packet count incremented	0
busy packet count decremented	0
session destructed in pme	0
session destruct set in pme	0
kq op hold	0
kq op drop	0
kq op route	0
kq op continue	0
kq op error	0
kq op stop	0
PME wait not set	0
PME wait set	0
PME KQ run not called	0

J-Web

- On all high-end SRX Series devices, on the Monitor > Events and Alarms > Security Events page, the *Is global policy* check box is introduced.
- On all branch SRX Series and J Series devices, the username field does not accept HTML tags or the "<" and ">" characters. The following error message appears:
A username cannot include certain characters, including < and >
- On all branch SRX Series devices, on the Monitoring Policies page, the Deactivate and Move functions on the toolbar and the Count and Log action columns in the output table are not supported and will no longer be available.
- On all branch SRX Series devices, on the Checking Policies page, the Delete and Deactivate buttons are not supported and will no longer be available.

Logical Systems

- In Junos OS releases earlier than Junos OS Release 12.1X46-D10, when a logical tunnel interface with an IPv4 address and an Ethernet encapsulation type is configured, a configuration check is performed to ensure that the address is not identical to its peer logical tunnel interface address and that both addresses are on the same subnet. However, when a logical tunnel interface with an IPv6 address and an Ethernet encapsulation type is configured, no such configuration check is performed.

Starting in Junos OS Release 12.1X46-D10, a check is performed for IPv6 configurations. However, this change can cause existing IPv6 configurations to fail.

Network Time Protocol

- When the NTP client or server is enabled in the **edit system ntp** hierarchy, the **REQ_MON_GETLIST** and **REQ_MON_GETLIST_1** control messages supported by the monlist feature within the NTP might allow remote attackers, causing a denial of service. To identify the attack, apply a firewall filter and configure the router's loopback address to allow only trusted addresses and networks.

Policy Applications

- In Junos OS releases earlier than Junos OS Release 12.1X46-D15, when you set the **count** option on a security policy using the CLI statement **security policies from-zone zone-name to-zone zone-name policy policy-name then**, the count is based on the number of packets and bytes of all network traffic that the policy allows to pass through the device.

In Junos OS Release 12.1X46-D15 and later, when you set the **count** option, the count is based on the number of packets and bytes of all network traffic the policy allows to pass through the device in both directions: the originating traffic from the client to the server (from the from-zone to the to-zone), and the return traffic from the server to the originating client.

Simple Network Management Protocol (SNMP)

- On all branch SRX Series and J Series devices, the screen SNMP trap **jnxJsScreenCfgChange** will not be sent during reboot.

System Logs

On all branch SRX Series devices, the following system log messages have been updated to include the **certificate ID**:

- PKID_PV_KEYPAIR_DEL
Existing message: **Key-Pair deletion failed**
New message: **Key-Pair deletion failed for <cert-id>**
- PKID_PV_CERT_DEL
Existing message: **Certificate deletion has occurred**

New message: **Certificate deletion has occurred for <cert-id>**

- PKID_PV_CERT_LOAD

Existing message: **Certificate has been successfully loaded**

New message: **Certificate <cert-id> has been successfully loaded**

- PKID_PV_KEYPAIR_GEN

Existing message: **Key-Pair has been generated**

New message: **Key-Pair has been generated for <cert-id>**

System Management

- During a load override, to enhance the memory for the commit script, make sure you load the configuration by applying the following commands before commit:
set system scripts commit max-datasize 800000000
set system scripts op max-datasize 800000000

User Interface and Configuration

- You can configure only one rewrite rule for one logical interface. When you configure multiple rewrite rules for one logical interface, an error message is displayed and the commit fails.

Virtual Private Networks (VPNs)

- In previous Junos OS releases, the Pulse client could be automatically downloaded and installed when users logged into a branch SRX Series device that was configured for dynamic VPN. Starting with Junos OS Release 12.1X46-D30, Pulse client software is no longer available from dynamic VPN SRX Series devices and must be obtained from the Juniper Networks Download Software site at <http://www.juniper.net/support/downloads/>
- On all branch SRX Series devices, for path MTU calculations, the IPsec authentication data length is fixed at 16 bytes. However, the authentication data length for packets going through the IPsec tunnel is in accordance with the authentication algorithm negotiated for that tunnel.

The authentication data lengths for the different algorithms are:

- hmac-md5-96 (12 bytes)
 - hmac-sha-256-128 (16 bytes)
 - hmac-sha1-96 (12 bytes)
- For each VPN tunnel, both ESP and AH tunnel sessions are installed on SPUs and the control plane. In previous Junos OS releases, two tunnel sessions of the same protocol (ESP or AH) were installed for each VPN tunnel. For branch SRX Series devices, tunnel sessions are updated with the negotiated protocol after negotiation is completed. For high-end SRX Series devices, tunnel sessions on anchor SPUs are updated with the negotiated protocol while non-anchor SPUs retain ESP and AH tunnel sessions.

The ESP and AH tunnel sessions are displayed in the outputs for the **show security flow session** and **show security flow cp-session** operational mode commands.

- As of Junos OS Release 11.4, checks are performed to validate the IKE ID received from the VPN peer device. By default, SRX Series and J Series devices validate the IKE ID received from the peer with the IP address configured for the IKE gateway. In certain network setups, the IKE ID received from the peer (which can be an IPv4 or IPv6 address, fully qualified domain name, distinguished name, or e-mail address) does not match the IKE gateway configured on the SRX Series or J Series device. This can lead to a Phase 1 validation failure.

To modify the configuration of the SRX Series or J Series device or the peer device for the IKE ID that is used:

1. On the SRX Series or J Series device, configure the **remote-identity** statement at the **[edit security ike gateway gateway-name]** hierarchy level to match the IKE ID that is received from the peer. Values can be an IPv4 or IPv6 address, fully qualified domain name, distinguished name, or e-mail address.



NOTE: If you do not configure **remote-identity**, the device uses the IPv4 or IPv6 address that corresponds to the remote peer by default.

2. On the peer device, ensure that the IKE ID is the same as the **remote-identity** configured on the SRX Series or J Series device. If the peer device is an SRX Series or J Series device, configure the **local-identity** statement at the **[edit security ike gateway gateway-name]** hierarchy level. Values can be an IPv4 or IPv6 address, fully qualified domain name, distinguished name, or e-mail address.
- The subject fields of a digital certificate can include Domain Component (DC), Common Name (CN), Organization Unit (OU), Organization (O), Location (L), State (ST), and Country (C).

In earlier releases, the **show security pki ca-certificate** and **show security pki local-certificate** CLI operational commands displayed only a single entry for each subject field, even if the certificate contained multiple entries for a field. For example, a certificate with two OU fields such as “OU=Shipping Department, OU=Priority Mail” displayed with only the first entry “OU=Shipping Department.” The **show security pki ca-certificate** and **show security pki local-certificate** CLI commands now display the entire contents of the subject field, including multiple field entries.

The commands also display a new subject string output field that shows the contents of the subject field as it appears in the certificate.

- When a remote user launches newly installed client software, the link to close the Web browser window does not appear in the VPN client launch page. The user must close the browser window by clicking the browser’s close button.
- Starting in Junos OS Release 12.1X46-D10, **local-address** can be configured at the **[edit security ike gateway gateway-name]** hierarchy level to specify the local gateway address when there are multiple addresses configured on an external physical interface to a VPN peer. **local-address** and the remote IKE gateway address must be in the same

address family, either IPv4 or IPv6. Prior to Junos OS Release 12.1X46-D10, **local-address** was a hidden CLI configuration statement.

Related Documentation

- [New and Changed Features on page 5](#)
- [Known Behavior on page 33](#)
- [Known Issues on page 53](#)
- [Resolved Issues on page 54](#)
- [Documentation Updates on page 76](#)
- [Migration, Upgrade, and Downgrade Instructions on page 86](#)

Known Behavior

Application Layer Gateways (ALGs)

- On all SRX Series devices, you can define the Sun RPC and MS-RPC mapping entry ageout value using the **set security alg sunrpc map-entry-timeout value** and **set security alg msrpc map-entry-timeout value** commands. The ageout value ranges from 8 hours to 72 hours, and the default value is 32 hours.

If either the Sun RPC ALG or the MS-RPC ALG service does not trigger the control negotiation even after 72 hours, the maximum RPC ALG mapping entry value times out and the new data connection to the service fails.

- The maximum size of the jbuf is 9 Kb. If the message buffer size is more than 9 Kb, the entire message cannot be transferred to the ALG packet handler. This causes subsequent packets in the session to bypass ALG handling, resulting in a transaction failure.

The limitations for SCCP ALGs are as follows:

- The SCCP is a Cisco proprietary protocol. So, any changes to the protocol by Cisco cause the SCCP ALG implementation to break. However, workarounds are provided to bypass strict decoding and allow any protocol changes to be handled gracefully.
- The SCCP ALG validates protocol data units (PDUs) with message IDs in the ranges [0x0 - 0x12], [0x20 - 0x49], and [0x81 - 0x14A]. By default, all other message IDs are treated as unknown messages and are dropped by the SCCP ALG.
- Any changes to the policies will drop the sessions and impact already established SCCP calls.
- The SCCP ALG opens pinholes that are collapsed during traffic or media inactivity. This means that during a temporary loss of connectivity, media sessions are not reestablished.
- CallManager (CM) version 6.x and later does not support TCP probe packets in chassis cluster mode. As a result, the existing SCCP sessions will break when there is a failover. You can still create new SCCP sessions during failover.

The PPTP ALG with IPv6 support has the following limitation:

- Because PPP packets are compressed with Microsoft Point-to-Point Encryption (MPPE) protocol after the tunnel is set up, translation of the IP header in the PPP package cannot be handled; therefore, to make sure PPTP connection works well, the PPTP client must be able to work in dual stack mode. So that an IPv6 PPTP client can accept an IPv4 address for PPP tunnel interface, by which it can communicate with the IPv4 PPTP server without IP address translation for PPP packets.

The RTSP ALG with IPv6 support has the following limitations:

- Real-Time Streaming Protocol (RTSP) is an Application Layer protocol for controlling the delivery of data with real-time properties. The RTSP ALG supports a peer client, and the server transmits real-time media; it does not support third-party endpoints involved in the transaction.
- In case of destination NAT or NAT64 for IP address translation, if the RTSP message (including the Session Description Protocol (SDP) application content) length exceeds 2500 bytes, then the RTSP ALG processes only the first 2500 bytes of the message and ignores the rest of the message. In this scenario, the IP address in the RTSP message is not translated if the IP address does not appear in the first 2500 bytes.

The SIP ALG with IPv6 support has the following limitation:

- When NAT64 with persistent NAT is implemented, the SIP ALG adds the NAT translation to the persistent NAT binding table if NAT is configured on the Address of Record (AOR). Because persistent NAT cannot duplicate the address configured, coexistence of NAT66 and NAT64 configured on the same address is not supported.

Only one binding is created for the same source IP address.

AppSecure

- J-Web pages for AppSecure are preliminary.
- Custom application signatures and custom nested application signatures are not currently supported by J-Web.
- When ALG is enabled, application identification includes the ALG result to identify the application of the control sessions. Application firewall permits ALG data sessions whenever control sessions are permitted. If the control session is denied, there will be no data sessions. When ALG is disabled, application identification relies on its signatures to identify the application of the control and data sessions. If a signature match is not found, the application is considered unknown. Application firewall handles applications based on the application identification result.

AX411 Access Points

- On SRX210, SRX240, and SRX650 devices, you can configure and manage a maximum of four access points.
- On all branch SRX Series devices, managing AX411 WLAN Access Points through a Layer 3 ae interface is not supported.

Chassis Cluster

- SRX100, SRX210, SRX240, and SRX650 devices have the following chassis cluster limitations:
 - VRRP is not supported.
 - Unified ISSU is not supported.
 - The 3G dialer interface is not supported.
 - On SRX Series device failover, access points on the Layer 2 switch reboot and all wireless clients lose connectivity for 4 to 6 minutes.
 - VDSL Mini-PIMs are not supported in chassis cluster.
 - Queuing on the ae interface is not supported.
 - Group VPN is not supported.
 - On SRX100 and SRX110 devices, switching is not supported in chassis cluster mode.
 - The Chassis Cluster MIB is not supported.
 - Any packet-based services such as MPLS and CLNS are not supported.
 - On the lsq-0/0/0 interface, Link services MLPPP, MLFR, and CRTP are not supported.
 - On the lt-0/0/0 interface, CoS for RPM is not supported.

Starting with Junos OS Release 12.1X45-D10 and later, sampling features such as flow monitoring, packet capture, and port mirroring are supported on reth interfaces.

- On all SRX Series devices in a chassis cluster, flow monitoring for version 5 and version 8 is supported. However, flow monitoring for version 9 is not supported.
- If you use packet capture on reth interfaces, two files are created, one for ingress packets and the other for egress packets based on the reth interface name. These files can be merged outside of the device using tools such as Wireshark or Mergecap.
- If you use port mirroring on reth interfaces, the reth interface cannot be configured as the output interface. You must use a physical interface as the output interface. If you configure the reth interface as an output interface using the **set forwarding-options port-mirroring family inet output** command, the following error message is displayed.

Port-mirroring configuration error.

Interface type in reth1.0 is not valid for port-mirroring or next-hop-group config

- Packet-based forwarding for MPLS and ISO protocol families is not supported.
- The factory default configuration for SRX100 devices automatically enables Layer 2 Ethernet switching. Layer 2 Ethernet switching is not supported in chassis cluster mode for SRX100 devices. If you use the factory default configuration, you must delete Ethernet switching before you enable chassis clustering.
- On all J Series devices, a Fast Ethernet port from a 4-port Ethernet PIM cannot be used as a fabric link port in a chassis cluster.

- On all branch SRX Series devices, reth interfaces and the lo0 interface are supported for IKE external interface configuration in IPsec VPN. Other interface types can be configured, but IPsec VPN might not work.
- On all J Series devices, the ISDN feature on chassis cluster is not supported.

Command-Line Interface (CLI)

- On all branch SRX Series and J Series devices, the **clear services flow** command is not supported.
- On all J Series devices, RADIUS accounting is not supported.
- On SRX210 and SRX240 devices, J-Web crashes if more than nine users log in to the device by using the CLI. The number of users allowed to access the device is limited as follows:
 - For SRX210 devices: four CLI users and three J-Web users
 - For SRX240 devices: six CLI users and five J-Web users
- On J6350 devices, there is a difference in the power ratings provided by user documentation (*J Series Services Routers Hardware Guide* and PIM, uPIM, and ePIM Power and Thermal Calculator) and the power ratings displayed by CLI (by a unit of 1). The CLI display rounds off the value to a lower integer, and the ratings provided in user documentation round off the value to the higher integer. As a workaround, follow the user documentation for accurate ratings.
- On all branch SRX Series devices, the tunnel-queuing option is not supported in chassis cluster mode.

Connectivity Fault Management (CFM)

- CFM is not supported on the following interfaces:
 - 8-Port Gigabit Ethernet SFP XPIM
 - 2-Port 10-Gigabit Ethernet XPIM
 - 1-Port SFP Mini-PIM
- CFM is supported only on interfaces with the Ethernet switching family.

Dynamic Host Configuration Protocol (DHCP)

- On all branch SRX Series devices, DHCP relay is unable to update the binding status based on DHCP_RENEW and DHCP_RELEASE messages.
- On all branch SRX Series and J Series devices, DHCPv6 client authentication is not supported.
- On all branch SRX Series and J Series devices, DHCP client and server functionality is not supported in a chassis cluster.
- On all branch SRX Series devices, DHCPv6 client does not support:

- Temporary addresses
- Reconfigure messages
- Multiple identity association for nontemporary addresses (IA_NA)
- Multiple prefixes in a single identity association for prefix delegation (IA_PD)
- Multiple prefixes in a single router advertisement

Flow and Processing

- On all branch SRX Series devices, GRE fragmentation is not supported in packet-based mode.
- On all branch SRX Series and J Series devices, a mismatch between the Firewall Counter Packet and Byte Statistics values, and between the Interface Packet and Byte Statistics values, might occur when the rate of traffic increases above certain rates of traffic.
- On SRX100, SRX210, SRX220, SRX240, and SRX650 devices, due to a limit on the number of large packet buffers, Routing Engine based sampling might run out of buffers for packet sizes greater than or equal to 1500 bytes and hence those packets will not be sampled. The Routing Engine could run out of buffers when the rate of the traffic stream is high.
- On SRX100 and SRX240 devices, the data file transfer rate for more than 20 Mbps is reduced by 60 percent with the introduction of Junos Pulse 1.0 client as compared to the Acadia client that was used before Junos OS Release 11.1.
- On SRX100, SRX210, SRX220, SRX240, and SRX650 devices, the default authentication table capacity is 10,000; the administrator can increase the capacity to a maximum of 15,000.
- On all branch SRX Series and J Series devices, when devices are operating in flow mode, the Routing Engine side cannot detect the path MTU of an IPv6 multicast address (with a large size packet).
- On all branch SRX Series devices, you cannot configure route policies and route patterns in the same dial plan.
- On all J Series devices, even when forwarding options are set to drop packets for the ISO protocol family, the device forms ES-IS adjacencies and transmits packets because ES-IS packets are Layer 2 terminating packets.
- On all branch SRX Series and J Series devices, high CPU utilization triggered for reasons such as CPU intensive commands and SNMP walks causes the BFD protocol to flap while processing large BGP updates.
- On SRX210, SRX240, and J Series devices, broadcast TFTP is not supported when flow is enabled on the device.

- On all branch SRX Series devices, the maximum number of concurrent sessions for SSH, Telnet, and Web is as follows:

Sessions	SRX100	SRX210	SRX220	SRX240	SRX550	SRX650
SSH	3	3	250	5	5	5
Telnet	3	3	250	5	5	5
Web	7	7	7	7	7	7



NOTE: These defaults are provided for performance reasons.

- On SRX210 and SRX240 devices, for optimized efficiency, we recommend that you limit use of CLI and J-Web to the numbers of sessions listed in the following table:

Device	CLI	J-Web	Console
SRX210	3	3	1
SRX240	5	5	1

- On SRX100 devices, Layer 3 control protocols (OSPF, using multicast destination MAC address) on the VLAN Layer 3 interface work only with access switch ports.

Hardware

- On all branch SRX Series devices, a chassis cluster is only supported when both devices are the same model and have the same amount of memory. Thus, a chassis cluster is not supported if it combines SRX Series branch devices with 1-GB and 2-GB memory in the same cluster.

Interfaces and Routing

- When using SRX Series devices in chassis cluster mode, we recommend that you do not configure any local interfaces (or combination of local interfaces) along with redundant Ethernet interfaces.

For example:

The following configuration of chassis cluster redundant Ethernet interfaces, in which interfaces are configured as local interfaces, is not recommended:

```

ge-2/0/2 {
  unit 0 {
    family inet {
      address 1.1.1.1/24;
    }
  }
}

```

The following configuration of chassis cluster redundant Ethernet interfaces, in which interfaces are configured as part of redundant Ethernet interfaces, is recommended:

```

interfaces {
  ge-2/0/2 {
    gigether-options {
      redundant-parent reth2;
    }
  }
  reth2 {
    redundant-ether-options {
      redundancy-group 1;
    }
    unit 0 {
      family inet {
        address 1.1.1.1/24;
      }
    }
  }
}

```

- On SRX100, SRX110, SRX210, and SRX220 devices, you cannot configure the same VRRP group ID on different interfaces of a single device.
- On all branch SRX Series devices, PIM does not support upstream and downstream interfaces across different virtual routers in flow mode
- On all branch SRX Series devices, the Link Layer Discovery Protocol (LLDP) is not supported on reth interfaces.
- On all J Series devices, the flow monitoring version 9 has the following limitations:
 - Routing Engine based flow monitoring V5 or V8 mode is mutually exclusive with inline flow monitoring V9.
 - Flow aggregation for V9 export is not supported.
 - Only UDP over IPv4 or IPv6 protocol can be used as the transport protocol.
 - Only the standard IPv4 or IPv6 template is supported for exporting flow monitoring records.
 - User-defined or special templates are not supported for exporting flow monitoring records.
- On all branch SRX Series and J Series devices, flow monitoring IPv6 version 9 has the following limitations:
 - MPLS in not supported.
 - User-defined version 9 templates are not supported.
 - Routing Engine based flow monitoring version 9 is not supported.
 - Flow monitoring and accounting are not supported in chassis cluster mode.
 - Flow monitoring and accounting are not supported on an ae interface.
 - J-Web for IPv6 sampled packets is not supported.

- SNMP queries for IPv6 sampled packets are not supported
- Flow monitoring can be configured in version 5, version 8, or version 9 export mode. Up to eight version 9 collectors are supported in export mode.
- Scope of accounting of IPv6 flow monitoring version 9 packets associated with pseudointerfaces (such as IRB, ML, LAG, VLAN, and GRE) is not supported.
- Creation of an SCTP session (parallel to TCP) between an exporter and a collector for gathering flow monitoring information is not supported.
- Maximum flow sessions that might be supported include:
 - A device with 1-GB RAM, such as an SRX220 device, might support up to 15,000 flow monitoring sessions at a time.
 - A device with 2-GB RAM, such as an SRX650 device, might support up to 59,900 flow monitoring sessions at a time.
- Changes in source AS and destination AS are not immediately reflected in exported flows.
- On all branch SRX Series devices, IPv6 traffic transiting over IPv4 based IP over IP tunnel (for example, IPv6-over-IPv4 using ip-x/x/x interface) is not supported.
- The ATM interface takes more than 5 minutes to come up when CPE is configured in ANSI-DMT mode and CO is configured in automode. This occurs only with ALU 7300 DSLAM, due to limitation in current firmware version running on the ADSL Mini-PIM.
- On SRX100 and J Series devices, dynamic VLAN assignments and guest VLANs are not supported.
- On all branch SRX Series devices, the subnet directed broadcast feature is not supported.
- On SRX650 devices, Ethernet switching is not supported on Gigabit Ethernet interfaces (ge-0/0/0 through ge-0/0/3 ports).
- On SRX210, SRX220, SRX240, and SRX650 devices, when using stream mode security logging, security logs cannot be sent to NSM or another syslog server if the server is in the same subnet as interface fxp0. Stream mode syslog can only be routed out via revenue ports and not via the fxp0 interface. This implies that you cannot configure the security log server in the same subnet as the fxp0 interface.
- On all branch SRX Series devices, the number of child interfaces per node is restricted to 4 on the reth interface and the number of child interfaces per reth interface is restricted to 8.
- On SRX240 High Memory devices, traffic might stop between the SRX240 device and the Cisco switch due to link mode mismatch. We recommend setting the same value to the autonegotiation parameters on both ends.
- On SRX100 devices, the link goes down when you upgrade FPGA on 1xGE SFP. As a workaround, run the **restart fpc** command and restart the FPC.
- On SRX210 devices with VDLS2, ATM COS VBR-related functionality cannot be tested.

- On SRX210 devices, IGMPv2 JOINS messages are dropped on an IRB interface. As a workaround, enable IGMP snooping to use IGMP over IRB interfaces.
- On all J Series devices, the DS3 interface does not have an option to configure multilink-frame-relay-uni-nni (MFR).
- On SRX210, SRX220, and SRX240 devices, every time the VDSL2 Mini-PIM is restarted in the ADSL mode, the first packet passing through the Mini-PIM is dropped.
- On all branch SRX Series devices, the RPM server operation does not work when the probe is configured with the option **destination-interface**.
- On all J Series devices, LLDP is not supported on routed ports.
- In J Series xDSL PIMs, mapping between IP CoS and ATM CoS is not supported. If the user configures IP CoS in conjunction with ATM CoS, the logical interface level shaper matching the ATM CoS rate must be configured to avoid congestion drops in segmentation and reassembly (SAR) as shown in the following example:


```
set interfaces at-5/0/0 unit 0 vci 1.110
set interfaces at-5/0/0 unit 0 shaping cbr 62400 ATM COS
set class-of-service interfaces at-5/0/0 unit 0 scheduler-map sche_map IP COS
set class-of-service interfaces at-5/0/0 unit 0 shaping-rate 62400 ADD IFL SHAPER
```
- On SRX650 devices, MAC pause frame and FCS error frame counters are not supported for the interfaces ge-0/0/0 through ge-0/0/3.
- On SRX240 and SRX650 devices, the VLAN range from 3967 to 4094 falls under the reserved VLAN address range, and the user is not allowed any configured VLANs from this range.
- On SRX650 devices, the last four ports of a 24-Gigabit Ethernet switch GPIM can be used either as RJ-45 or small form-factor pluggable transceiver (SFP) ports. If both are present and providing power, the SFP media is preferred. If the SFP media is removed or the link is brought down, then the interface will switch to the RJ-45 medium. This can take up to 15 seconds, during which the LED for the RJ-45 port might go on and off intermittently. Similarly, when the RJ-45 medium is active and an SFP link is brought up, the interface will transition to the SFP medium, and this transition could also take a few seconds.
- On SRX210 devices, the USB modem interface can handle bidirectional traffic of up to 19 Kbps. On oversubscription of this amount (that is, bidirectional traffic of 20 Kbps or above), keepalives do not get exchanged, and the interface goes down.
- On SRX100, SRX210, SRX240, and SRX650 devices, on the Layer 3 ae interface, the following features are not supported:
 - Encapsulations (such as CCC, VLAN CCC, VPLS, and PPPoE)
 - J-Web
 - 10-Gigabit Ethernet
- On SRX100 devices, the multicast data traffic is not supported on IRB interfaces.
- On SRX240 High Memory devices, when the **system login deny-sources** statement is used to restrict the access, it blocks a remote copy between nodes, which is used to

copy the configuration during the commit routine. Use a firewall filter on the lo0.0 interface to restrict the Routing Engine access. However, if you choose to use the **system login deny-sources** statement, check the private addresses that were automatically on lo0.x and sp-0/0/0.x and exclude them from the denied list.

- On SRX100, SRX210, SRX220, SRX240, SRX650, and all J Series devices, on VLAN-tagged routed interfaces, LLDP is not supported.
- On SRX210 devices, the DOCSIS Mini-PIM delivers speeds up to a maximum of 100 Mbps throughput in each direction.
- On SRX550 and SRX650 devices, the aggregate Ethernet (ae) interface with XE member interface cannot be configured with the Ethernet switching family.
- On all branch SRX Series and J Series devices, the Q-in-Q support on a Layer 3 interface has the following limitations:
 - Double tagging is not supported on reth and ae interfaces.
 - Multitopology routing is not supported in flow mode and in chassis clusters.
 - Dual tagged frames are not supported on encapsulations (such as CCC, TCC, VPLS, and PPPoE).
 - On Layer 3 logical interfaces, input-vlan-map, output-vlan-map, inner-range, and inner-list are not applicable
 - Only TPIDs with 0x8100 are supported, and the maximum number of tags is 2.
 - Dual tagged frames are accepted only for logical interfaces with IPV4 and IPV6 families.
- On SRX650 devices, LLDP is not supported on the base ports of the device and on the 2-Port 10 Gigabit Ethernet XPIM.
- On SRX100, SRX110, SRX210, SRX220, SRX240, and SRX550 devices, LACP is not supported on the 1-Port Gigabit Ethernet SFP Mini-PIM.
- On all branch SRX Series devices, IKEv2 does not include support for:
 - Policy-based tunnels
 - Dial-up tunnels
 - NAT-T
 - VPN monitoring
 - NHTB for st0—Reusing the same tunnel interface for multiple tunnels
 - EAP
 - Multiple child SAs for the same traffic selectors for each QoS value
 - Proposal enhancement features
 - Reuse of DH exponentials
 - Configuration payloads

- IP Payload Compression Protocol (IPComp)
- Dynamic endpoint (DEP) VPN

Intrusion Detection and Prevention (IDP)

- On all branch SRX Series devices, from Junos OS Release 11.2 and later, the IDP security package is based on the Berkeley database. Hence, when the Junos OS image is upgraded from Junos OS Release 11.1 or earlier to Junos OS Release 11.2 or later, a migration of IDP security package files needs to be performed. This is done automatically on upgrade when the IDP process comes up. Similarly, when the image is downgraded, a migration (secDb install) is automatically performed when the IDP process comes up, and previously installed database files are deleted.

However, migration is dependent on the XML files for the installed database present on the device. For first-time installation, completely updated XML files are required. If the last update on the device was an incremental update, migration might fail. In such a case, you have to manually download and install the IDP security package using the **download** or **install** CLI command before using the IDP configuration with predefined attacks or groups.

As a workaround, use the following CLI commands to manually download the individual components of the security package from the Juniper Security Engineering portal and install the full update:

- **request security idp security-package download full-update**
- **request security idp security-package install**
- On all branch SRX Series devices, IDP does not allow header checks for nonpacket contexts.
- On SRX100, SRX210, SRX220, SRX240, and SRX650 devices, the maximum supported number of entries in the ASC table is 100,000 entries. Because the user land buffer has a fixed size of 1 MB as a limitation, the table displays a maximum of 38,837 cache entries.
- On all branch SRX Series devices, with regard to serialization limits, the maximum number of IDP sessions supported is shown in [Table 4 on page 43](#):

Table 4: Maximum Number of IDP Sessions

Branch SRX Series Device	1-GB Memory	2-GB Memory
SRX100 and SRX110	16,000	16,000
SRX210	16,000	32,000
SRX220	16,000	32,000
SRX240	32,000	64,000
SRX550	32,000	64,000
SRX650	32,000	64,000

- On all branch SRX Series devices, all IDP policy templates are supported except All Attacks. There is a 100 MB policy size limit for integrated mode and a 150 MB policy size limit for dedicated mode. The current supported IDP policy templates are dynamic based on the attack signatures added. Therefore, be aware that supported templates might eventually grow past the policy size limit.

On all branch SRX Series devices, the following IDP policies are supported:

- DMZ_Services
 - DNS_Service
 - File_Server
 - Getting_Started
 - IDP_Default
 - Recommended
 - Web_Server
- On all branch SRX Series devices, IDP deployed in both active/active and active/passive chassis clusters has the following limitations:
 - No inspection of sessions that fail over or fail back.
 - The IP action table is not synchronized across nodes.
 - The Routing Engine on the secondary node might not be able to reach networks that are reachable only through a Packet Forwarding Engine.
 - The SSL session ID cache is not synchronized across nodes. If an SSL session reuses a session ID and it happens to be processed on a node other than the one on which the session ID is cached, the SSL session cannot be decrypted and will be bypassed for IDP inspection.
 - On all branch SRX Series devices, IDP deployed in active/active chassis clusters has a limitation that for time-binding scope source traffic, if attacks from a source (with more than one destination) have active sessions distributed across nodes, then the attack might not be detected because time-binding counting has a local-node-only view. Detecting this sort of attack requires an RTO synchronization of the time-binding state that is not currently supported.



NOTE: On SRX100 devices, IDP chassis cluster is supported in active/backup mode.

IPv6

- **Network and Security Manager (NSM)**—Consult the NSM release notes for version compatibility, required schema updates, platform limitations, and other specific details regarding NSM support for IPv6 addressing on SRX Series and J Series devices.

J-Web

- **SRX Series and J Series browser compatibility**
 - To access the J-Web interface, your management device requires the following software:
 - Language support—English-version browsers
 - Supported OS—Microsoft Windows XP Service Pack 3
 - Supported browsers

Device	Application	Supported Browsers	Recommended Browser
SRX100, SRX110, SRX210, SRX220, SRX240, SRX550, SRX650	J-Web	<ul style="list-style-type: none"> • Mozilla Firefox version 3.x • Microsoft Internet Explorer version 7.0 <p>NOTE: The New Setup wizard and the PPPoE wizard work best with Mozilla Firefox version 15.x or later.</p>	Mozilla Firefox version 3.x

- To use the Chassis View, a recent version of Adobe Flash that supports ActionScript and AJAX (Version 9) must be installed. Also note that the Chassis View is displayed by default on the Dashboard page. You can enable or disable it using options in the Dashboard Preference dialog box, but clearing cookies in Microsoft Internet Explorer also causes the Chassis View to be displayed.
- On all branch SRX Series devices, in the J-Web interface, there is no support for changing the T1 interface to an E1 interface or vice versa. As a workaround, use the CLI to convert from T1 to E1 and vice versa.
- On all branch SRX Series and J Series devices, users cannot differentiate between Active and Inactive configurations on the System Identity, Management Access, User Management, and Date & Time pages.
- On SRX210 devices, there is no maximum length when the user commits the hostname in CLI mode; however, only 58 characters, maximum, are displayed in the J-Web System Identification panel.
- On all J Series devices, some J-Web pages for new features (for example, the Quick Configuration page for the switching features on J Series devices) display content in one or more modal pop-up windows. In the modal pop-up windows, you can interact only with the content in the window and not with the rest of the J-Web page. As a result, online Help is not available when modal pop-up windows are displayed. You

can access the online Help for a feature only by clicking the Help button on a J-Web page.

- On all branch SRX Series devices, you cannot use J-Web to configure a VLAN interface for an IKE gateway. VLAN interfaces are not currently supported for use as IKE external interfaces.

The PPPoE wizard has the following limitations:

- While you use the load and save functionality, the port details are not saved in the client file.
- The Non Wizard connection option cannot be edited or deleted through the wizard. Use the CLI to edit or delete the connections.
- The PPPoE wizard cannot be launched if the backend file is corrupted.
- The PPPoE wizard cannot be loaded from the client file if non-wizard connections share the same units.
- The PPPoE wizard cannot load the saved file from one platform to another platform.
- There is no backward compatibility between PPPoE wizard Phase 2 to PPPoE wizard Phase 1. As a result, the PPPoE connection from Phase 2 will not be shown in Phase 1 when you downgrade to an earlier release.

The New Setup wizard has the following limitations:

- The Existing Edit mode might not work as expected if you previously configured the device manually, without using the wizard.
- Edit mode might overwrite outside configurations such as Custom Application, Policy Name, and zone inbound services.
- In create new mode, when you commit your configuration changes, your changes will overwrite the existing configuration.
- VPN and NAT wizards are not compatible with the New Setup wizard; therefore the VPN or NAT wizard configuration will not be reflected in the New Setup wizard or vice versa.
- By default, 2 minutes are required to commit a configuration using the New Setup wizard.
- On SRX650 devices, the default mode configures only the ge-0/0/1 interface under the internal zone.
- You might encounter usability issues if you use Microsoft Internet Explorer version 7 or 8 to launch the New Setup wizard.
- If you refresh your browser after you download the license, the factory mode wizard is not available.
- When you commit the configuration, the underlying Web management interface changes, and you do not receive a response about the commit status.
- Webserver ports 80 (HTTP) and 443 (HTTPS) on the DMZ or internal zone are overshadowed if Web management is enabled on the Internet zone not configured for

destination NAT. As a workaround, change the webserver port numbers for HTTP and HTTPS by editing the recommended policies on the Security policies page.

- Images, buttons, and spinner (indicating that the configuration is being applied) on the wizard screen do not initially appear when the browser cache is cleared.

Layer 2 Transparent Mode

- DHCP server propagation is not supported in Layer 2 transparent mode.

Network Address Translation (NAT)

- **Single IP address in a source NAT pool without PAT**—The number of hosts that a source NAT pool without PAT can support is limited to the number of addresses in the pool. When you have a pool with a single IP address, only one host can be supported, and traffic from other hosts is blocked because there are no resources available.

If a single IP address is configured for a source NAT pool without PAT when NAT resource assignment is not in active-backup mode in a chassis cluster, traffic through node 1 will be blocked.

- For all ALG traffic, except FTP, we recommend that you not use the static NAT rule options **source-address** or **source-port**. Data session creation can fail if these options are used, because the IP address and the source port value, which is a random value, might not match the static NAT rule. For the same reason, we also recommend that you not use the source NAT rule option **source-port** for ALG traffic.

For FTP ALG traffic, the **source-address** option can be used because an IP address can be provided to match the source address of a static NAT rule.

Additionally, because static NAT rules do not support overlapping addresses and ports, they should not be used to map one external IP address to multiple internal IP addresses for ALG traffic. For example, if different sites want to access two different FTP servers, the internal FTP servers should be mapped to two different external IP addresses.

- Maximum capacities for source pools and IP addresses have been extended on SRX650 devices, as follows:

Devices	Source NAT Pools	PAT Maximum Address Capacity	Pat Port Number	Source NAT Rules Number
SRX650 (High Memory devices)	1024	1024	64M	1024
SRX650 (Low Memory devices)	256	256	16M	1024

Increasing the capacity of source NAT pools consumes memory needed for port allocation. When source NAT pool and IP address limits are reached, port ranges should be reassigned. That is, the number of ports for each IP address should be decreased when the number of IP addresses and source NAT pools is increased. This ensures NAT

does not consume too much memory. Use the **port-range** statement in configuration mode in the CLI to assign a new port range or the **pool-default-port-range** statement to override the specified default.

Configuring port overloading should also be done carefully when source NAT pools are increased.

For source pool with PAT in range (63,488 through 65,535), two ports are allocated at one time for RTP/RTCP applications, such as SIP, H.323, and RTSP. In these scenarios, each IP address supports PAT, occupying 2048 ports (63,488 through 65,535) for ALG module use.

- **NAT rule capacity change**—To support the use of large-scale NAT at the edge of the carrier network, the device-wide NAT rule capacity has been changed.

The number of destination and static NAT rules has been incremented as shown in [Table 5 on page 48](#). The limitation on the number of destination-rule-set and static-rule-set has been increased.

[Table 5 on page 48](#) provides the requirements per device to increase the configuration limitation as well as to scale the capacity for each device.

Table 5: Number of Rules on SRX Series and J Series Devices

NAT Rule Type	SRX100	SRX210	SRX240	SRX650	J Series
Source NAT rule	512	512	1024	1024	512
Destination NAT rule	512	512	1024	1024	512
Static NAT rule	512	512	1024	6144	512

The restriction on the number of rules per rule set has been increased so that there is only a device-wide limitation on how many rules a device can support. This restriction is provided to help you better plan and configure the NAT rules for the device.

- On all branch SRX Series devices, in case of SSL proxy, sessions are whitelisted based on the actual IP address and not on the translated IP address. Because of this, in the whitelist configuration of the SSL proxy profile, the actual IP address should be provided and not the translated IP addresses.

Example:

Consider a destination NAT rule that translates destination IP address 20.20.20.20 to 5.0.0.1 using the following commands:

- **set security nat destination pool d1 address 5.0.0.1/32**
- **set security nat destination rule-set dst-nat rule r1 match destination-address 20.20.20.20/32**
- **set security nat destination rule-set dst-nat rule r1 then destination-nat pool d1**

In the above scenario, to exempt a session from SSL proxy inspection, the following IP address should be added to the whitelist:

- **set security address-book global address ssl-proxy-exempted-addr 20.20.20.20/32**
- **set services ssl proxy profile ssl-inspect-profile whitelist ssl-proxy-exempted-addr**

Power over Ethernet (PoE)

- On SRX210-PoE devices, SDK packages might not work.

Security Policies

- On all branch SRX Series devices, the current SSL proxy implementation has the following connectivity limitations:
 - The SSLv2 protocol is not supported. SSL sessions using SSLv2 are dropped.
 - SSL sessions where client certificate authentication is mandatory are dropped.
 - SSL sessions where renegotiation is requested are dropped.
- On all branch SRX Series devices, for a particular session, the SSL proxy is only enabled if a relevant feature related to SSL traffic is also enabled. Features that are related to SSL traffic are IDP, application identification, application firewall, and application tracking. If none of the above listed features are active on a session, the SSL proxy bypasses the session and logs are not generated in this scenario.
- On all branch SRX Series and J Series devices, you cannot configure the following IP addresses as negated addresses in a policy:
 - Wildcard addresses
 - IPv6 addresses
 - Addresses such as any, any-ipv4, any-IPv6, and 0.0.0.0
- When a range of addresses or a single address is negated, it can be divided into multiple addresses. These negated addresses are shown as a prefix or a length that requires more memory for storage on a Packet Forwarding Engine.
- Each platform has a limited number of policies with negated addresses. A policy can contain 10 source or destination addresses. The capacity of the policy depends on the maximum number of policies that the platform supports.
- J Series devices do not support the authentication order **password radius** or **password ldap** in the **edit access profile *profile-name* authentication-order** command. Instead, use **order radius password** or **ldap password**.

Simple Network Management Protocol (SNMP)

- On all J Series devices, the SNMP NAT related MIB is not supported.

Switching

- **Layer 2 transparent mode support**—On SRX100, SRX110, SRX210, SRX220, SRX240, SRX550, and SRX650 devices, the following features are not supported for Layer 2 transparent mode:

- G-ARP on the Layer 2 interface
- STP
- IP address monitoring on any interface
- Transit traffic through IRB
- IRB interface in a routing instance
- IRB interface handling of Layer 3 traffic



NOTE: The IRB interface is a pseudointerface and does not belong to the reth interface and redundancy group.

- On SRX100, SRX210, SRX240, and SRX650 devices, change of authorization is not supported with 802.1x.
- On SRX100, SRX110, SRX210, SRX240, SRX550, and SRX650 devices, on the routed VLAN interface, the following features are not supported:
 - IPv6 (family inet6)
 - IS-IS (family ISO)
 - Class of service
 - Encapsulations (Ether CCC, VLAN CCC, VPLS, PPPoE, and so on) on VLAN interfaces
 - CLNS
 - PIM
 - DVMRP
 - VLAN interface MAC change
 - G-ARP
 - Change VLAN-Id for VLAN interface

Unified Access Control

- During SRX device communication to the Infranet Controller (IC), the connection remains in attempt-next state preventing a successful communication. This happens when an outgoing interface used to connect the IC is a part of routing-instance.

Unified Threat Management (UTM)

- The quarantine action is supported only for UTM Enhanced Web Filtering or Juniper enhanced type of Web filtering.
- On SRX550 devices configured with Sophos Antivirus, certain files whose sizes are larger than the max-content-size might not go into fallback unlike other AV engines and instead end up being detected as clean file for few protocols which does not pre-declare the content size.

Upgrade and Downgrade

- On all J Series devices, the Junos OS upgrade might fail due to insufficient disk space if the CompactFlash is smaller than 1 GB in size. We recommend using a 1-GB compact flash for Junos OS Release 10.0 and later.
- On SRX100, SRX210, SRX220, SRX240, and SRX650 devices, when you connect a client running Junos Pulse 1.0 to an SRX Series device that is running a later version of Junos Pulse, the client will not be upgraded automatically to the later version. You must uninstall Junos Pulse 1.0 from the client and then download the later version of Junos Pulse from the SRX Series device.
- On the SRX240B2 and SRX240H2 models, when you try to upgrade from Junos OS Release 11.4 to Junos OS Release 12.1X44, 12.1X45, 12.1X46, or 12.1X47, the upgrade fails when attempting to validate the configuration. To resolve this, use the **no-validate** option.

USB

- On all branch SRX Series devices, frequent plug and play of USB keys is not supported. You must wait for the device node creation before removing the USB key.

Virtual Private Networks (VPNs)

- On all branch SRX Series devices, when you download the Pulse client using the Mozilla browser, the “Launching the VPN Client” page is displayed when Junos Pulse is still downloading. However, when you download the Pulse client using Microsoft Internet Explorer, the “Launching the VPN Client” page is displayed after Junos Pulse has been downloaded and installed.
- On SRX100, SRX210, SRX240, and SRX650 devices, while configuring dynamic VPN using the Junos Pulse client, when you select the authentication-algorithm as sha-256 in the IKE proposal, the IPsec session might not get established.

- RIP is not supported in point-to-multipoint (P2MP) VPN scenarios including AutoVPN deployments. We recommend OSPF or IBGP for dynamic routing when using P2MP VPN tunnels.
- On all branch SRX Series devices, configuring XAuth with AutoVPN secure tunnel (st0) interfaces in point-to-multipoint mode and dynamic IKE gateways is not supported.

The IPv6 IPsec VPN implementation has the following limitations:

- Devices with IPv6 addressing do not perform fragmentation. IPv6 hosts should either perform path MTU discovery or send packets smaller than the IPv6 minimum MTU size of 1280 bytes.
- Because IPv6 addresses are 128 bits long compared to IPv4 addresses, which are 32-bits long, IPv6 IPsec packet processing requires more resources. Therefore, a small performance degradation is observed.
- The dynamic VPN server must be a standalone branch SRX Series device. The dynamic VPN feature is not supported on high-end SRX Series devices or on branch SRX Series devices in a chassis cluster.
- The IPv6 IPsec VPN does not support the following functions:
 - Remote Access—XAuth, config mode, and shared IKE identity with mandatory XAuth
 - IKE authentication—PKI or DSA
 - IKE peer type—Dynamic IP
 - NAT-T
 - VPN monitoring
 - NHTB
 - Packet reordering for IPv6 fragments over tunnels is not supported
 - IPv6 link-local address

See *VPN Feature Support for IPv6 Addresses* for more information about IPv6 address support in VPN features.

On all branch SRX Series devices, when you enable VPN, overlapping of the IP addresses across virtual routers is supported with following limitations:

- An IKE external interface address cannot overlap with any other virtual router.
- An internal/trust interface address can overlap across virtual routers.
- An st0 interface address cannot overlap in route-based VPN in point-to-multipoint tunnels such as NHTB.
- An st0 interface address can overlap in route-based VPN in point-to-point tunnels.

SRX100, SRX210, and SRX240 devices have the following limitations:

- The IKE configuration for the Junos Pulse client does not support the hexadecimal preshared key.

- The Junos Pulse client IPsec does not support the AH protocol and the ESP protocol with NULL authentication.
- When you log in through the Web browser (instead of logging in through the Junos Pulse client) and a new client is available, you are prompted for a client upgrade even if the **force-upgrade** option is configured. Conversely, if you log in using the Junos Pulse client with the **force-upgrade** option configured, the client upgrade occurs automatically (without a prompt).

Related Documentation

- [New and Changed Features on page 5](#)
- [Changes in Behavior and Syntax on page 17](#)
- [Known Issues on page 53](#)
- [Resolved Issues on page 54](#)
- [Documentation Updates on page 76](#)
- [Migration, Upgrade, and Downgrade Instructions on page 86](#)

Known Issues

The following problems currently exist in Juniper Networks branch SRX Series Services Gateways and J Series Services Routers. The identifier following the description is the tracking number in the Juniper Networks Problem Report (PR) tracking system.



NOTE: For the latest, most complete information about outstanding and resolved issues with the Junos OS software, see the Juniper Networks online software defect search application at <http://www.juniper.net/prsearch>.

Interfaces and Routing

- On all branch SRX Series devices, packet loss is seen initially for 15 seconds after the GRE tunnel is brought up over the VDSL interface. [PR821330](#)
- On SRX550 and SRX650 devices, 20 to 40 percent traffic loss is seen on the port of the SRX-GP-2XE-SFP-PTX after changing the speed from 10 GB to 1 GB. This issue is seen in both fiber and copper mode. When you switch between fiber and copper mode on the port of the SRX-GP-2XE-SFP-PTX, the speed might vary within the configuration. [PR1033369](#)

Intrusion Detection and Prevention (IDP)

- On all branch SRX Series devices, an active IDP session might not be fetched by the **show security flow session idp** and **show security flow session summary idp** commands. This issue occurs only when both IDP and AppID are enabled. [PR1045587](#)

Network Address Translation (NAT)

- On all branch SRX Series devices, when source NAT is configured, the ports are allocated randomly by default. In rare circumstances, the global random port table of source pools or interfaces becomes damaged by certain services or traffic. This damage can result in low-range ports being assigned a higher priority in sessions. Ports might be reused quickly, causing application access failure. [PR1006649](#)

Platform and Infrastructure

- On all branch SRX Series devices, when you connect to the device through wireless AP the secure access port incorrectly allows access to the MAC addresses that are not in the list of allowed MAC addresses. [PR587163](#)

Related Documentation

- [New and Changed Features on page 5](#)
- [Changes in Behavior and Syntax on page 17](#)
- [Known Behavior on page 33](#)
- [Resolved Issues on page 54](#)
- [Documentation Updates on page 76](#)
- [Migration, Upgrade, and Downgrade Instructions on page 86](#)

Resolved Issues

The following are the issues that have been resolved in Junos OS Release 12.1X46 for Juniper Networks SRX Series Services Gateways. The identifier following the description is the tracking number in the Juniper Networks Problem Report (PR) tracking system.



NOTE: For the latest, most complete information about outstanding and resolved issues with the Junos OS software, see the Juniper Networks online software defect search application at <http://www.juniper.net/prsearch>.

Resolved Issues - 12.1X46-D35

Application Identification

- On all branch SRX Series devices running Junos OS Release 12.1X46 and earlier, if application identification (AppID) is enabled, performance degradation is seen in comparison with devices running Junos OS Release 12.1X47-D10 and later. This is because the AppID function does not ignore the related sessions when AppID has reached the terminal state, and continues with the serialization processing for those sessions. It is important to note that Junos OS Release 12.1X47 and later releases use advanced AppID. [PR1046509](#)

Application Layer Gateways (ALGs)

- On all branch SRX Series devices (except SRX110) in a chassis cluster with TCP-based ALG enabled and the TCP keepalive mechanism used on the TCP server and client, after a data plane Redundancy Group (RG1+) failover, the keep-alive message causes the mbuf to be held by the ALG until the session timeout. As a result, a high mbuf usage alarm is generated. Application communication failure occurs due to lack of mbuf. [PR1031910](#)
- On all branch SRX Series devices with the SIP ALG and NAT enabled, if you place a call on hold or off hold many times, each time with different media ports, the resource in the call is used, resulting in one-way audio. Tearing down the call clears the resource, and following calls are not affected. [PR1032528](#)
- On all branch SRX Series devices (except SRX110) in a chassis cluster with the SCCP ALG enabled and if the SCCP state in use flag is not configured in the process of the SCCP call in the device, the related real-time object (RTO) hot synchronization might cause the flowd process to crash. [PR1034722](#)
- On all branch SRX Series devices with the MS-RPC ALG enabled, the flowd process might crash when the MS-RPC ALG processes the crafted **ISystemActivator RemoteCreateInstance Response** packets. [PR1036574](#)
- On all branch SRX Series devices with the SIP ALG and NAT enabled, the SIP ALG does not execute IP translation for the retransmitted 183 session progress messages. In this scenario, the SIP call will fail when the device receives the first 183 session progress messages without SDP information, but the retransmitted 183 session progress messages contains SDP information. [PR1036650](#)
- On all branch SRX Series devices, the DNS ALG does not terminate the session when a truncated DNS reply is received. Hence, the session remains up until high timeout (10~50) is reached. [PR1038800](#)
- On all branch SRX Series devices, SIP ALG code has been enhanced to support RFC 4566 regarding the SDP lines order and to avoid issues of no NAT in owner filed (O line) in some circumstances. [PR1049469](#)
- On all branch SRX Series devices with the MS-RPC ALG enabled, the flowd process might crash due to incorrect MS-RPC ALG parsing for the **ISystemActivator RemoteCreateInstance Response** packets. [PR1066697](#)

Authentication

- On all branch SRX Series devices with firewall authentication enabled, when a firewall authentication from an authenticated IP address for a new authentication fails, and then a pass-through firewall authentication tries this entry, the firewall authentication function accesses a freed memory, which results in a flowd process crash. [PR1040214](#)

Chassis Cluster

- On all branch SRX Series devices in chassis cluster mode, during control plane RGO failover, a policy resynchronization operation compares the policy message between the Routing Engine and the Packet Forwarding Engine. However, some fields in the security policy data message are not processed. Data for unprocessed fields might be treated differently and cause the flowd process to crash. [PR1040819](#)
- On all branch SRX Series devices in a chassis cluster, if the switching fabric (swfab) interface is configured, the swfab interface incorrectly updates the state of the fabric (fab) interface. As a result, the fab interface might be stuck in the down state. [PR1064005](#)

Dynamic Host Configuration Protocol (DHCP)

- On all branch SRX Series devices configured as a DHCP server (using the jdhcpd process), when the DHCP server gets a new request from a client and applies an IP address from the authentication process (authd), the jdhcpd process communicates with authd process twice as expected (once for the DHCP discovery message and once for the DHCP request message). If the authentication fails in the first message, the authd process will indefinitely wait for the second authentication request. However, the jdhcpd process never sends the second request, because the process detects that the first authentication did not occur. This causes memory leak on the authd process, and the memory might get exhausted, generating a core file and preventing DHCP server service. High CPU usage on the Routing Engine might also be observed. [PR1042818](#)

Flow-Based and Packet-Based Processing

- On all branch SRX Series devices, when composite next hop is used, RSVP session flap might cause an if state mismatch between the master Routing Engine and the backup Routing Engine, leading to a kernel crash on the master Routing Engine. [PR905317](#)
- On all branch SRX Series devices with IDP configured, in rare cases, where the device runs out of memory, the flowd process might crash if shell code detection occurs. [PR985139](#)
- On all branch SRX Series devices, when you configure **http-get RPM probes** to measure the website response, the probes might fail because the HTTP server might incorrectly interpret the request coming from the device. [PR1001813](#)
- On all branch SRX Series devices, IPsec tunnel reconnection might cause a memory leak. [PR1002738](#)
- On all multiple thread-based branch SRX Series devices (SRX240, SRX550, and SRX650), if IDP, AppSecure, ALG, GTP, or the SCTP feature, which is required for serialization flow processing is enabled, the device might encounter an issue where

two flow threads work on the same session at the same time for the serialization flow processing. This issue might cause memory corruption, and then result in a flowd process crash. [PR1026692](#)

- On all branch SRX Series devices, when you enable **flexible-vlan-tagging**, the return traffic might be dropped on the tagged interface with the following message: **packet dropped, pak dropped due to invalid l2 broadcast/multicast addr.** [PR1034602](#)
- On all branch SRX Series devices in a chassis cluster Z mode, if static NAT or destination NAT is configured, and in the NAT rule, the IP address of the incoming interface is used as a matching condition of the destination address (for example, **set security nat static <rule-set-name> rule <rule-name> match destination-address <use the IP address of incoming interface>**), then the traffic matching the NAT rule is discarded. [PR1040185](#)
- On all branch SRX Series devices with GRE tunnel configured, the carrier interface of GRE tunnel is not updated when a more accurate and new route to the tunnel destination address is added, which might cause traffic loss in some scenarios. [PR1040666](#)
- On all branch SRX Series devices, after **IDP drop** action is performed on a TCP session, the TCP session timeout is not accurate. [PR1052744](#)
- On all branch SRX Series devices running Junos OS Release 12.3X48-D10 or later, with enhanced Web filtering configured, the connection to the Websense ThreatSeeker Intelligence Cloud might time out if **strict-syn-check** is enabled under the **[security flow tcp-session]** hierarchy. [PR1061064](#)
- On SRX550 devices, traffic processed by the serialization process is dropped when the maximum limit of serialization sessions (32,000) is exceeded. As a result, advanced services such as IDP, ALG, and AppSecure are impacted. [PR1061524](#)

Hardware

- On all branch SRX Series devices, the message **twsi0: Device timeout on unit 1** fills the console on soft reboot. [PR1050215](#)

Interfaces and Routing

- On all branch SRX Series devices, the **clear security dns-cache** command is extended to resolve all DNS entries immediately. Similarly, the security policies containing DNS names are updated immediately to use the refreshed IP addresses after the FQDN addresses are resolved. [PR970235](#)
- On SRX100H2, SRX110H2, SRX210H2, SRX220H2, and SRX240H2 devices, when you enable vlan tagging on interfaces and commit the configuration, the interface speed and duplex mode might cause the interface to stop processing traffic. [PR1003423](#)
- On SRX240, SRX550, and SRX650 devices, a delay of several seconds (maximum 4 seconds) might occur to detect that the link is down. [PR1008324](#)
- On all branch SRX Series devices configured as a CHAP authentication client, in a PPPoE over ATM LLC encapsulation scenario, the connection might not be established because of an incorrect sequence of messages being exchanged with the second LNS. [PR1027305](#)

- On all branch SRX Series devices, the **commit synchronize** command fails because the kernel socket gets stuck. [PR1027898](#)
- On all branch SRX Series devices, multiple CoS rewrite rules are applied to a single interface where only one rewrite rule is allowed. [PR1034173](#)
- On all branch SRX Series devices in a chassis cluster with PPPoE configured on a redundant Ethernet (reth) interface, when both nodes reboot, the PPPoE interface (pp0.x) sometimes is not prepared, despite the PPPoE session being up. [PR1050264](#)
- On all branch SRX Series devices with PPPoE configured, when PPPoE fails to authenticate, the software next-hop entry will leak in the data plane, gradually consuming all 64,000 software next-hop entries. When the software next-hop table is full, the following next-hop error pops up: **RT_PFE: NH IPC op 2 (CHANGE NEXTHOP) failed, err 6 (No Memory) peer_class 0, peer_index 0 peer_type 10**. [PR1055882](#)
- On all branch SRX Series devices, the very-high-bit-rate digital subscriber line (VDSL) firmware upgrade fails due to a permission issue, and the error message **No applicable firmware present** is displayed. [PR1066032](#)

Installation and Upgrade

- On SRX650 devices, if the u-boot revision is 2.5 or later, installing the Junos OS release image from TFTP in loader mode fails. [PR1016954](#)

Intrusion Detection and Prevention (IDP)

- On SRX210 and SRX220 devices, due to memory constraints, the combination of large IDP policies (that is, IDP_Default) along with express antivirus (EAV) might not compile successfully. [PR974851](#)
- On all branch SRX Series devices, when IDP and Express Antivirus (EAV) are configured under very high stress, application traffic might coredump. [PR1019401](#)

J-Web

- On all branch SRX Series devices, J-Web sets a limitation on the size of the configuration fetched from a device to avoid memory exhaustion. When the configuration size exceeds this limitation, J-Web fails to load the configuration on Junos OS Release 12.3X48-D10. [PR1037073](#)
- On all branch SRX Series devices, J-Web does not display all the member link interfaces for aggregate Ethernet (ae) interface. [PR1038850](#)
- On all branch SRX Series devices, security policy log or security policy count is not displayed when the match condition is **RT_FLOW_SESSION**. [PR1056947](#)
- On all branch SRX Series devices, when you use a configuration encryption, the missing rescue configuration alarm is set even when there is a saved rescue configuration. [PR1057473](#)
- On all branch SRX Series devices, if a security policy contains a **tcp-options** statement, modifying this security policy by using J-Web results in the loss of the **tcp-options** statement. This is because the **tcp-options** configuration is missing in the J-Web security policy configuration. [PR1063593](#)

Network Address Translation (NAT)

- On all branch SRX Series devices with persistent NAT enabled, if an invalid flow with the protocol value 0 creates a persistent NAT entry, then this persistent NAT entry is not cleared even when the invalid session is cleared. [PR935325](#)

Platform and Infrastructure

- On all branch SRX Series devices, after enabling IEEE 802.1X, the connected devices on some ports might fail to be authenticated. This is because MAC authentication requests might get stuck on the eswd process, therefore this issue might be seen on certain random ports, not all ports. [PR1042294](#)
- On all branch SRX Series devices, the configurations of group junos-defaults are lost after a configuration roll back. As a result, the **commit** command fails. [PR1052925](#)
- On SRX100 devices, when you run the **show snmp mib walk jnxMibs** command, the chassisd log repeatedly generates the **fru is present: out of range slot -1 for FAN** message. [PR1062406](#)
- On all branch SRX Series devices, the log displays the message **log: /kernel: veriexec: fingerprint for dev**. This is a cosmetic issue. [PR1064166](#)
- On SRX100 devices, when the device is configured as an authentication enforcer of 802.1x, authentication from certain special supplicants might fail. This is because the software engine that processes the next-hops in the device incorrectly processes the packet coming from the supplicant with a special source MAC address. As a result, the packets are dropped. [PR1067588](#)

Security Policy

- On all branch SRX Series devices, when two security policies are combined and the whole address space is used, then the secondary security policy might fail to evaluate traffic. [PR1052426](#)
- On all branch SRX Series devices, changing a dynamic address of a security policy might cause its dynamic address identification to be mismatched between the Routing Engine and the Packet Forwarding Engine due to the difference between the new and the old configuration being ignored. [PR1061253](#)

Unified Threat Management (UTM)

- On all branch SRX Series devices, when UTM Sophos antivirus is enabled and a file that is not supported by Sophos antivirus is transferred through SMTP, the device might not be able to handle the last packet, and mail will be on hold. When packets are later sent on this session, the packet that was on hold will be handled by the device and the system will return to normal state. [PR1049506](#)
- On all branch SRX Series devices, if the name server is configured and the interface pointing to the name server is down, in a rare condition, the flowd process might crash due to a UTM internal function even though UTM is not configured. [PR1066510](#)

Virtual Private Networks (VPNs)

- On SRX240, SRX550, and SRX650 devices with IPsec VPN configured using IKE version 1, the device can hold only two pairs of IPsec security associations (SAs) per tunnel. When the third IPsec SA rekey occurs, the oldest IPsec SA is deleted. Due to this mechanism, a looping of IPsec SA rekey might occur. For example, when a VPN peer contains incorrect configuration that has more than two proxy IDs matching only one proxy ID on a device, the rekey looping issue might cause the flowd process to crash on multiple thread-based SRX Series devices. [PR996429](#)
- On all branch SRX Series devices, in a hub-and-spoke IPsec VPN scenario, on the hub site, when committing the static NHTBs on the multipoint secure tunnel (st0) interface, the VPN routes might become active even though the VPN tunnel is down. This issue also occurs when the system reboots with static NHTBs and the related static routes configured. [PR1007235](#)
- On all branch SRX Series devices, in group VPN setups, all the already registered members might suddenly disappear from the key server due to memory leak. [PR1023940](#)
- On all branch SRX Series devices, when IPsec VPN is enabled using IKE version 2 and a distinguished name is used to verify the IKE version 2 Phase 1 remote identity, a remote peer initiates IKE version 2 Phase 1 security association (SA) renegotiation (SRX Series devices work as responders), the new negotiated VPN tunnel might stay in "inactive" state on the data plane, causing IPsec VPN traffic loss. [PR1028949](#)
- On all branch SRX Series devices, the block size for Advanced Encryption Standard (AES) in Galois/Counter Mode (GCM) has been reduced from 8 to 4. Block size 8 is used for connecting to other SRX Series devices, and block size 4 is interoperable with systems from Cisco, strongSwan, and other companies. When you set the correct block size 4 for AES-GCM, it causes a problem when connecting to previous releases of Junos OS for SRX Series devices. The problem affects certain packet sizes, so it might appear to work for some traffic, such as ping, but not for other traffic. In a hub-and-spoke configuration, the upgrade causes problems with tunnels to all spokes until they are upgraded. [PR1037432](#)
- On all branch SRX Series devices, when a primary IP address of an interface changes, some IPsec tunnels terminated on that interface might go down. [PR1044620](#)
- On all branch SRX Series devices configured with a large number of IPsec VPN tunnels, in a very rare condition, if VPN monitoring is enabled, the kmd process might crash when you delete the partial VPN tunnels. [PR1044660](#)

Resolved Issues - 12.1X46-D30

Application Layer Gateways (ALGs)

- On all branch SRX Series devices with SIP ALG enabled, when either retain-hold-resource and NAT are configured or retransmission of 183 session progress messages with SDP occurs (the first transmission did not have SDP), the SIP ALG incorrectly changes the IP address that is embedded inside the media payload to zero, causing a call failure. [PR1016969](#)

- On all branch SRX Series devices, in certain situations, the H.323 ALG incorrectly handles translation because the stored position is not initialized properly. As a result, H.323 endpoints registration failure and call failure occur. [PR1023528](#)

Chassis Cluster

- On all branch Series devices in a chassis cluster, the security zone is not populated properly on the J-Web interface port configuration page. [PR859200](#)
- On all branch SRX Series devices, in dual fabric link chassis clusters, when the control link and one fabric link go down, the chassis cluster goes into a split brain condition in which both nodes become primary. With one fabric link up, the secondary node of the chassis cluster goes into an ineligible state and then into the disabled state. [PR989548](#)
- On SRX100, SRX110, and SRX210 devices, no events are displayed when the temperature of the chassis exceeds the thermal threshold value. [PR999888](#)
- On all branch SRX Series devices configured in a chassis cluster, VLAN interfaces on the primary node might flap or become down. [PR1001162](#)
- On all branch SRX Series devices in a chassis cluster, when the “switch to L2 mode” button is pressed in J-Web, it does not ask for any confirmation and converts to transparent mode immediately and reboots the device. [PR1007740](#)

Command Line Interface (CLI)

- On all branch SRX Series devices, the Network Security Daemon (NSD) process might crash, causing the **show security match-policies** command to generate multiple core files. This event occurs because, the policy database does not synchronize between the Routine Engine and the Packet Forwarding Engine. [PR1003099](#)
- On all branch SRX Series devices, the CLI auto-complete does not work for any key words after **set system login class <name> permissions** command. [PR1032498](#)

Dynamic Host Configuration Protocol (DHCP)

- On all branch SRX Series devices, if the DHCPv6 client is configured for the PPPoE interface and the pp0 interface is disabled and again enabled, the pp0 interface does not acquire the IPv6 address from the DHCPv6 server. [PR998712](#)
- On all branch SRX Series devices, in DHCP requests, the IP TTL value is set to 1 and the DHCP option 12 is missing. [PR1011406](#)
- When a branch SRX Series device is configured as a DHCP server (using JDHCP), inspite of explicitly configuring **next-server** (siaddr) and **tftp boot-server** option, the siaddr and tftp boot server are set as 0.0.0.0 in DHCP reply packets. [PR1034735](#)

Flow and Processing

- On all branch SRX Series devices, under certain conditions, the creation of a multicast leaf session might result in an invalid multicast next hop, which crashes the flow module. [PR921438](#)
- On all branch SRX Series devices, the temporary flowd process crashes while you run the **get-software-information level=detail** command using a NETCONF client. This type of flowd crash is harmless. [PR937450](#)
- On all branch SRX Series devices, CoS buffer sizes are not recalculated after you delete the interface units. This might result in suboptimal CoS behavior. [PR953924](#)
- On SRX240, SRX550, and SRX650 devices, in multithreaded, mixed traffic (TCP or UDP) environments, packets might go out of order or get dropped by the device. [PR977614](#)
- On all branch SRX Series devices, for IDP, AppSecure, ALG, GTP, or SCTP, the flow serialization impacts session performance. This flow serialization continues even after Layer 7 processing is completed. [PR986326](#)
- On all branch SRX Series devices in Layer 2 transparent mode, the flowd process might crash when two packets with the same connections are received in a short time before the flow session is created, and destination MAC address lookup succeeds for these two packets. [PR1025983](#)

Interfaces and Routing

- On SRX220 and SRX550 devices, you can configure a maximum of 250 connections as connection-limit. However, 250 connections cannot be established. To set the maximum-connection-limit, use the **set system services telnet connection-limit** command. [PR976318](#)
- On all branch SRX Series devices, when the **packet-capture** option is configured on the egress interface and a multicast stream is sent through the device, the multicast traffic might not be captured. [PR1005116](#)
- On all branch SRX Series devices, when a new user is created, the home directory for the user is not created. [PR1015156](#)
- On all branch SRX Series devices, in a rare condition, during a failure of routing update, a free memory might be accessed again, which results in the flowd process crash. [PR1017148](#)
- On all branch SRX Series devices, the flowd process might crash while applying a CoS filter for the host outbound traffic. [PR1021150](#)
- On all branch SRX Series devices with First Hop Router (FHR) in multicast scenario, after the device reboots, the PIM tunnel selects loopback0.0 as the outgoing interface due to a timing issue where the route is not ready. If the loopback0.0 and the downstream interface are not in the same security zone, the PIM register packets will be dropped because of reroute failure. [PR1031185](#)

Intrusion Detection and Prevention (IDP)

- On SRX240H2, SRX240H2-POE, and SRX240H2-DC devices, the IDP cannot process any traffic due to incorrect setting of flow sessions. [PR1011057](#)

J-Web

- On all branch SRX Series devices, on the Dashboard page, the serial number and the system uptime are not displayed. [PR1009371](#)
- On all branch SRX Series devices, the PKI certificate issued by J-Web GUI HTTPs will not be used when DVPN is configured in the same device. This is because the device will use the self-signed PKI certificate for both J-Web GUI HTTPs and DVPN URL access. [PR1017747](#)

Network Address Translation (NAT)

- On J Series devices, multicast traffic is not forwarded if source NAT is used on the traffic. [PR782159](#)

Platform and Infrastructure

- On SRX100B, SRX210B, and SRX240B devices, high control plane memory usage is expected after you upgrade to Junos OS Release 12.1X45 or 12.1X46. [PR985479](#)

Security

- OpenSSL released a Security Advisory that included CVE-2014-3566 known as the "POODLE" vulnerability. The SSL protocol 3.0 (SSLv3) uses nondeterministic CBC padding, which makes it easier for man-in-the-middle attackers to obtain clear text data through a padding oracle attack. OpenSSL is upgraded to support for SSL 3.0 fallback protection (TLS_FALLBACK_SCSV). Refer to JSA10656 for more information. [PR1033938](#)

System Log

- On all branch SRX Series devices, the "flowd_octeon_hm: pconn_client_connect: Failed to connect to the server after 0 retries" message repeats in the log. [PR1035936](#)

Unified Threat Management (UTM)

- On all branch SRX Series devices with the UTM Kaspersky antivirus (KAV) feature enabled, the chunked HTTP traffic might be terminated unexpectedly by the client due to incorrect content sent by the SRX Series devices. As a result, the whole page or partial content is not displayed in the client browser. [PR971895](#)
- On all branch SRX Series devices with UTM content filtering enabled, when the filename extension value is set to .com to block the URLs, the content filtering feature incorrectly treats the <searchpart> as a path and blocks the URLs that ends with .com. [PR1008108](#)

VPN

- On all branch SRX Series devices, the certificate-based IKEv2 tunnel might not be established if the wildcard (*) is configured as a remote identity for the IKE gateway. [PR968614](#)
- On all branch SRX Series devices, dynamic VPN user groups are not able to access certain remote resources. In this scenario, there are two policies referring to the same dynamic VPN and one of the policy directions is not set. Hence, the lookup fails in the null policy at the end. [PR988263](#)
- On all branch SRX Series devices, in group VPN member, the **KMD_PM_IKE_SERVER_NOT_FOUND** message appears repeatedly in the kmd log after rekey. [PR991306](#)
- On all branch SRX Series devices deployed in a hub-and-spoke VPN scenario with a hub point with dynamic endpoint VPN (DEP VPN) spokes, if manual NHTBs are configured, changing (adding or deleting) the NHTBs might cause other NHTBs to be deleted and existing tunnels to go down. [PR1001692](#)
- On all branch SRX Series devices in a Dynamic End Point (DEP) VPN scenario, the VPN tunnel might stay in down state after you change the user-at-hostname value. [PR1029687](#)
- On all branch SRX Series devices, in AutoVPN configuration after reboot, the VPN tunnel might not come up and an error with the private key is reported. [PR1032840](#)

Resolved Issues - 12.1X46-D25

Application-Aware Quality of Service (AppQoS)

- On all branch SRX Series devices, application traffic control rate limiters are not supported on model H2. [PR979901](#)

Dynamic Host Configuration Protocol (DHCP)

- On all branch SRX Series devices, when the DHCP client (a windows PC) only sends one DISCOVER packet, the DHCP server (an SRX Series device) receives two DISCOVER packets and replies with two OFFER packets. However, although it is not a problem to allocate the IP address of the DHCP client. [PR894760](#)

Flow and Processing

- On SRX Series devices, multicast traffic might cause memory leak on the data plane. [PR947894](#)
- On all branch SRX Series devices, the G-ARP replies do not update the existing MAC address entry. When the MAC address timer expires, a new MAC address is updated. [PR953879](#)
- On SRX240, SRX550, and SRX650 devices, in certain situations, flow sessions time out and get corrupted. This leads to the flow sessions being set to an abnormally high value, which eventually leads to the session table becoming full. [PR955630](#)

- On all branch SRX Series devices, the packets through IPsec VPN fail in chassis cluster Z mode when there is a fragmentation required. [PR956808](#)
- On all branch SRX Series devices deployed in a multicast scenario, a memory leak on the fwdd process might occur when the multicast routes change. [PR963116](#)
- On all branch SRX Series devices, the GRE tunnel does not change the outbound interface when the route changes. [PR965890](#)
- On all branch SRX Series devices with selective stateless packet-based services configured, self-traffic generated on custom routing instances will be dropped if it is forwarded in packet-based mode. [PR968631](#)
- On SRX550 devices, the maximum flow sessions are configured incorrectly. The devices have larger session capacities than the configured session values. [PR977169](#)
- On all branch SRX Series devices, due to an indirect next-hop change, memory corruption occurs in the flow route lookup table, which causes the flowd process to crash. [PR988659](#)

Interfaces and Routing

- On all branch SRX Series devices with 3G wireless modems, the 3G dialer interface dl0.0 might get stuck in the down link state. [PR855897](#)
- On all branch SRX Series devices, when you configure an ICMP **probe-server** option under the [**services rpm**] hierarchy for a specific interface (for example, ge-0/0/0), the device does not respond to ICMP requests from this interface. Other interfaces are not affected and continue to respond to ICMP requests. [PR960932](#)
- On SRX650 devices, the VLAN interface is down after a reboot due to a timing issue. [PR969079](#)
- On SRX550 and SRX650 devices with WAN cards installed, if an interface is configured for Ethernet switching mode and forwarding traffic, traffic processing might exhaust the mbuf pool. As a result, an interprocess communication (IPC) issue can occur, causing the WAN cards to go offline randomly. . [PR972332](#)

Intrusion Detection and Prevention (IDP)

- On all branch SRX Series devices, when you disable the **idp policy-optimizer** option using the **set security idp sensor-configuration no-policy-optimizer** command, the policy fails to load after reboot. [PR883258](#)

J-Web

- In J-Web, the App-FW page does not show the counter information. [PR972473](#)
- On all branch SRX Series devices, when you open several connections to J-Web from the same IP address, the HTTP process might hang and J-Web becomes unresponsive. [PR974042](#)

Network Address Translation (NAT)

- On all branch SRX Series devices, when the proxy-ndp feature is enabled on the interface, the entries in the IPv6 neighbor table from the interface might flap. [PR970281](#)

Platform and Infrastructure

- When all branch SRX Series devices are configured to use RADIUS authentication, and if the user-permission string sent from the RADIUS server is longer than 129 characters, the devices fail to process this user-permission string. This results in user permissions not being set correctly. [PR736331](#)

Switching

- On SRX210HE devices, after reboot, sometimes the VLAN interface is down while its member physical interface is up. [PR791610](#)

System Log

- On SRX650 devices, when you execute the **show security nat static rule all** command continuously, the following message is displayed:

kern.maxfiles limit exceeded by uid 0

[PR721715](#)

- On all branch SRX Series devices, every time a user logs in with SSH, a **verifexec: fingerprint mismatch** message is reported in the log. [PR929612](#)

Unified Threat Management (UTM)

- On branch SRX Series devices with UTM Sophos antivirus (SAV) service enabled, if source NAT for self-generated traffic is configured, the DNS queries from the UTM SAV service fail as timeout. [PR963978](#)

Virtual Private Networks (VPNs)

- On all branch SRX Series devices, in a hub-and-spoke IPsec VPN scenario, on the hub site, when you commit the static NHTBs on the multipoint secure tunnel (st0) interface, the VPN routes might become active even though the VPN tunnel is down. This issue also occurs when you reboot the system with static NHTBs and the related static routes configured. [PR947149](#)
- On all branch SRX Series devices, IPsec VPN tunnels could not come up due to unavailability of buffer space. [PR985494](#)

Resolved Issues - 12.1X46-D20

Chassis Cluster

- On all branch SRX Series devices in a chassis cluster, the counter for incoming traffic on a fabric interface always shows zero (0). [PR949962](#)
- On all branch SRX Series devices (except the SRX110), in an asymmetric chassis cluster, the secondary node (node1) uses a local interface to back up the interface in the primary

node (for example, node 0). If there is a route change, then the traffic is sent to the egress from the backup interface, which is the local interface of node 1. After the route resumes, the traffic is sent back to the egress from the primary interface, which is the local interface of node 0. The session related to the route change is in active state on both the nodes. Traffic might be interrupted when the session times out on the backup node and the session on the primary node is deleted. [PR951607](#)

Flow and Processing

- On SRX240, SRX550, and SRX650 devices, when the device receives a TCP reset (RST) and a FIN (the second FIN of the session) at the same time for a session, the RST and the FIN packet might get processed by different threads. As a result, the session timeout updates incorrectly, and the session remains on the session table for 150 seconds. [PR950799](#)
- On all branch SRX Series devices in a site-to-site VPN scenario, when the device is configured as an IPsec initiator, the flow session timeout is refreshed by the reroute packet. This causes an old session to remain in the session table, the VPN connection not to recover, and packet drops to occur. [PR959559](#)
- On all branch SRX Series devices with the IP spoofing screen enabled, the routing table search fails when it is locked by the system. As a result, false positives occur on IP spoofing detection. [PR967406](#)

Hardware

- On SRX100, SRX210, and SRX240 model B and H devices with 1 GB of RAM, the predefined IPS templates other than the recommended template might not compile successfully. [PR925337](#)

Interfaces and Routing

- On all branch SRX Series devices, because of a timing issue, the VLAN interface might fail to add security zone information after the RGO failover. [PR944017](#)
- On all branch SRX Series devices with interfaces encapsulated with ethernet-ccc, when you connect to an ae interface with Link Aggregation Control Protocol (LACP) enabled, the LACP packets do not pass through the ethernet-ccc encapsulated interface. [PR945004](#)
- On all branch SRX Series devices, when RGO failover is triggered, the old RGO primary device reboots or sometimes both the devices reboot. [PR953723](#)
- On SRX100B2, SRX100H2, SRX210B, SRX210HE2, SRX210HE2POE, SRX220H2, SRX220H2POE, SRX240B, SRX240B2, SRX240H2, and SRX240H2POE devices, the PPPoE feature session is disconnected or the connection is not available. [PR956307](#)

J-Web

- When you change the password minimum-length characters from 6 to 8, J-Web shows the error message **minimum-length is 6**. [PR942219](#)
- On all branch SRX Series devices, J-Web does not display the log sessions. [PR962892](#)

Platform and Infrastructure

- On all branch SRX Series devices, when using JDHCP, the server does not respond to the client with the DHCPOFFER packet when it receives the DHCPDISCOVER packet from the client. This causes the authd process to consume a large amount of CPU usage and increases the /mfs partition storage capacity. [PR925111](#)
- On all branch SRX Series devices, SSH connection is not possible between Cisco devices running IOS version 15 or later and SRX Series devices running Junos OS Release 11.2 or later. [PR957483](#)
- On J Series devices, kernel warnings about **kern.maxproc** nearing the limit value might appear in the log. [PR958358](#)

System Log

- On all branch SRX Series devices, the following error message is displayed on system or event logs after you upgrade to Junos OS Release 12.1X46-D10: **Can't find ifa on e1-x/0/x.y**. This message is harmless and does not affect the E1 interfaces and can be ignored. [PR971503](#)

Unified Threat Management (UTM)

- On all branch SRX Series devices, the **test security utm anti-virus** command for the antivirus feature does not work. [PR951124](#)

Virtual Private Networks (VPNs)

- Certificate-based authentication would fail when the RSA signature from the remote peer used SHA-256 as the message digest algorithm. [PR936141](#)
- On all branch SRX Series devices configured as a route-based IPsec Dynamic End Point (DEP) VPN node, the VPN tunnel interface st0.X link incorrectly remains up when IPsec Security Association (SA) is not established, even though VPN monitoring or **establish-tunnels immediately** is configured. [PR947552](#)
- On all SRX Series devices, in some situations, if the CRL server is not reachable, a memory leak might occur and show the message **kern.maxfiles limit exceeded by uid 0** in the console mode. Hence, the device administrator is not able to log in to the device anymore. [PR959194](#)
- On all branch SRX Series devices, when dynamic VPN is configured, it is not possible to use **local-certificate** or **pki-local-certificate** for Web management. A commit error is displayed when these options are configured. Only the self-signed certificate option can be configured. [PR969672](#)

Resolved Issues - 12.1X46-D15

Application-Aware Quality of Service (AppQoS)

- When GRE is enabled, AppQoS classification, marking, or rate limit does not work for fragmented packets in the client-to-server direction. [PR924932](#)

Dynamic Host Configuration Protocol (DHCP)

- In the DHCPv6 client command description, the word *stateful* was misspelled as *statefull*. It is changed to *stateful* in the description; however, the keyword is retained as *stateful* to avoid incompatibility. [PR924692](#)

Flow and Processing

- On SRX240, SRX550, and SRX650 devices, when the device receives out-of-order packets while transferring large TCP files, the throughput might be heavily impacted. [PR881761](#)
- On devices with 1 GB of memory, if the advanced services license is configured with the **reduce-dp-memory** option, memory is not released from the data plane to the control plane. [PR895648](#)
- On all SRX Series devices, if GRE tunnel configuration is committed without a correct route to the tunnel destination, the GRE tunnel session will bind the wrong anchor interface (the GRE tunnel outgoing interface) by route lookup. This anchor interface will not be updated even after the route is corrected when you commit the subsequent configuration. [PR933591](#)
- On all branch SRX devices, when the device is in packet mode, after you change an interface configuration, the warning message "warning: You have changed inet flow mode; You must reboot the system for your change to take effect" is displayed. The same message is displayed on every commit until the next reboot. This message can be safely ignored. [PR949472](#)
- On SRX210 devices running in packet mode, when DSCP marking (32 - 63) is on and the destination MAC in the packet header is present in the SRX ARP table, the devices reply to packets that are not destined to them. On devices in a chassis cluster, you must ensure that packets not destined to the SRX210 do not reach the device. [PR950486](#)

Hardware

- On the B and H versions of SRX100, SRX210, and SRX240 devices with 1 GB of RAM, the predefined IPS templates other than the recommended template might not compile successfully. [PR925337](#)
- On SRX550 and SRX650 devices, the SRX-GP-DUAL/QUAD-T1-E1 GPIM might have interoperability issues with the remote CSU using national standard feature due to the violation of ITU-T recommendation G.704. [PR939944](#)

Interfaces and Routing

- On SRX550 devices, the T3/E3 FPC goes offline after provisioning a switched port. [PR919617](#)
- On SRX Series devices with the 3G USB wireless modem, when the signal is low, the 3G cellular modem interface (cl-0/0/*) displays the status as Connected even though there is no signal or there is a low signal with no network connection. This is because there is no mechanism for the wireless WAN process to notify the Routing Engine status change even though the Packet Forwarding Engine is notified. After the signal recovers, the 3G cellular modem interface is not able to dial again. [PR923056](#)
- On SRX550 devices with DS3/E3 interfaces, the external clocking option is disabled to support the clocking option. [PR936356](#)

Screens

- When you use the **screen ids-option limit-session destination-ip-based** command, the session synchronization is not correct. [PR940029](#)

Unified Threat Management (UTM)

- On all branch SRX Series devices with the UTM Kaspersky antivirus (KAV) option enabled, and the intelligent-prescreening option configured, the chunked packet that only contains chunk-size data without any actual data is recognized as an invalid data packet, and the packet is dropped before it passes to the KAV engine in the KAV HTTP proxy processing. [PR937539](#)

Virtual Private Networks (VPNs)

- Certificate-based authentication would fail when the RSA signature from the remote peer used SHA-256 as the message digest algorithm. [PR936141](#)
- On all SRX Series devices, when IPsec is used in a chassis cluster, after the SPU or flowd uptime reaches 50 days or more, the amount of RTO traffic on the fabric link increases. [PR941999](#)
- On devices in a chassis cluster, during RGO failover to new primary node, if a route-based VPN does not have IPsec SAs associated with the tunnel, then the bind interface (st0) associated with the tunnel is marked as down. [PR944478](#)
- After the traffic-selector configuration is deleted from the VPN configuration object, the data traffic stops passing through the tunnel. [PR944598](#)

Resolved Issues - 12.1X46-D10

Application Layer Gateways (ALGs)

- The total SIP call values were incorrect, and the ALG feature could not be verified. [PR839190](#)
- On all branch SRX Series devices in a chassis cluster, the flowd process might crash when the ALG is enabled and the security policy is configured with the **log** option for ALG traffic. [PR889097](#)

- The Sun RPC ALG cannot open the gate as expected if the port string in **get-address** message is longer than 6, because current Sun RPC ALG can only parse the **uaddr** port string which is lesser than 6. [PR901205](#)

Authentication

- On all branch SRX Series devices configured with firewall authentication, if a user was already authenticated, and then when a subsequent user initiated authentication using the same IP address as the first user, the subsequent user inherited the first authenticated user's "Access time remaining" value. [PR843591](#)

BGP

- Under specific time-sensitive circumstances, if BGP determines that an UPDATE is too big to be sent to a peer, and immediately attempts to send a withdraw message, the routing daemon (rpd) may crash. An example of an oversized BGP UPDATE is one where a very long AS_PATH would cause the packet to exceed the maximum BGP message size (4096 bytes). The use of a very large number of BGP Communities can also be used to exceed the maximum BGP message size.

Please refer to JSA10609 for additional information. [PR918734](#)

Chassis Cluster

- On devices in a chassis cluster, when you execute the **clear system commit** command, it clears commit only from the local node. [PR821957](#)
- On devices in a chassis cluster, during a control link failure, if the secondary node is rebooted by control link failure recovery, the rebooted node goes into disabled state even after startup. [PR828558](#)
- On all branch SRX Series devices in a chassis cluster, when you download IDP Signature Database from the primary node, the **sig-db** is not synchronized to the secondary node. [PR914987](#)

Command-Line Interface (CLI)

- There is no specific CLI command to display the count of sessions allowed, denied, or terminated because of UAC enforcement. [PR733995](#)
- AppQoS does not display the correct application identification name when you run the **show class-of-service application-traffic-control statistics rate-limiter** CLI command. [PR751490](#)
- Certain combinations of Junos OS CLI commands and arguments have been found to be exploitable in a way that can allow root access to the operating system. This may allow any user with permissions to run these CLI commands the ability to achieve elevated privileges and gain complete control of the device.

Please refer to JSA10608 for additional information. [PR912707](#), [PR913328](#), [PR913449](#), [PR913831](#), [PR915313](#), [PR915957](#), [PR915961](#), [PR921219](#), [PR921499](#)

Dynamic Host Configuration Protocol (DHCP)

- On all branch SRX Series devices, when there are multiple interfaces configured as DHCP client, if one of DHCP client interface is from down state to up state, the IP address acquired by other DHCP client interfaces will be deleted unexpectedly and are added back after sometime. There is temporary traffic interruption until the deleted IP address is recovered automatically. [PR890124](#)
- Prior to Junos OS Release 11.4R9, DHCP option 125 cannot be configured for use as the **byte-stream** option. With Junos OS Release 11.4R9 and later releases, DHCP option 125 can be used for the **byte-stream** option. [PR895055](#)
- On all branch SRX Series devices working as DHCP clients, when the connection with the primary DHCP server is lost, the device tries to renew the lease. The device then drops the DHCP rebind ACK from the other DHCP server, which tries to assign the same IP address to it. [PR911864](#)

Flow and Processing

- When DNS ALG was enabled, the rewrite rules applied on the egress interface might not work for DNS messages. [PR785099](#)
- After enabling IPv6 in flow mode, IPv6 routes are not active. [PR824563](#)
- Current implementation of timeout for http is 1800s, the default timeout should be 300s. [PR858621](#)
- The RPM script triggers twice when the RPM probe-test fails. [PR869519](#)
- On J Series devices, the self-originating outbound traffic always uses the first logical unit queue. [PR887283](#)
- On all branch SRX Series devices with the MS-RPC ALG enabled, when the **junos-ms-rpc** application is not configured in the security policy and if the MS RPC control session is permitted by the security policy that matched the application “any”, then the MS-RPC ALG should not check the MS RPC data session and be permitted by the security policy. If the MS-RPC data session is configured to be processed by one or more other services such as IDP, UTM, AppID, or AppFW, then the MS-RPC ALG incorrectly checks the MS RPC data session and discards the MS RPC data session. [PR904682](#)
- On SRX100, SRX110, SRX210, and SRX220 devices with the FTP Application Layer Gateway (ALG) enabled, ICMP redirect might not work for FTP traffic. [PR904686](#)
- The memory allocated for multicast session might not release when multicast reroute occurs, this leads to memory leak. [PR905375](#)
- When you use a classifier based on EXP bits on a PE router, the CoS marked MPLS traffic is forwarded to the default egress queues instead of the custom configured queues. [PR920066](#)

Hardware

- On SRX210 devices, after you upgrade to Junos OS 12.1X46-D10 or later, the fan speed in relation to the Routing Engine temperature does not follow the temperature threshold table. [PR910977](#)

Infrastructure

- On all branch SRX Series devices, when the device authentication is through RADIUS server and the password protocol is Microsoft CHAP version 2, the password change operation fails as the user password change is enforced through Microsoft Active Directory server. [PR740869](#)
- After an upgrade, you cannot copy files between nodes in a cluster using the **file copy** command. [PR817228](#)
- On SRX240 devices, when a nonstandard HTTPS port is set, the URI changes to the IP address and port. [PR851741](#)
- On SRX100B and SRX100H devices unexpected system reboot is observed, and multiple core files are generated due to a double data rate2 (DDR2) memory timing issue between DRAM and CPU. The symptoms include flowd core files, core files from other daemons (such as snmpd, ntpd, rtlogd and so on) and silent reboot without core files are generated. These core files are related to random memory access (example: pointer corruption in session ager ring entry). [PR909069](#)

Interfaces and Routing

- The Routing Protocol Daemon (RPD) might crash with the following error: **/kernel: BAD_PAGE_FAULT: pid 1472 (rpd), uid 0: pc 0x86ff81c got a read fault at 0x15,x86 fault flags = 0x4**, when the OSPF switches from the primary path to the secondary path when loop-free alternates (LFA) and LDP-SYNC are enabled. The corruption is caused when OSPF does not completely free a memory location that is later reused by LDP. [PR737141](#)
- On VLAN tagged ethernet frames (802.1p), you cannot modify the VDSL priority bits. [PR817939](#)
- On SRX550 devices, the VRRP does not work when it is connected through IRB. [PR834766](#)
- On J Series devices, a Layer 2 loop might occur for a short time when you run the **request system power-off**, **request system reboot**, or **request system halt** command. [PR856457](#)
- The RPM script is triggered twice when the RPM probe-test fails. [PR869519](#)
- When a SHDSL Mini-PIM is configured in 2-wire mode with annex mode as Annex B/G, one of the physical interfaces does not come up. [PR874249](#), [PR882035](#)
- The point-to-multipoint (P2MP) interface does not accept any multicast packets, this leads to interoperability issues with the Secure Services Gateway (SSG). [PR895090](#)
- When there is a configuration change in the VDSL profile from one to another, the VDSL line does not retrain and comes up with the newly configured VDSL profile. [PR898775](#)

- When the virtual routers (routing instances) are connected with a looped cable and if one of the interfaces is VLAN, the unicast communication is unsuccessful. [PR909190](#)
- When multiple routing-instances are defined, DNS names in the address-book entries might not get resolved. This results in corresponding security policies to be nonoperational. [PR919810](#)

J-Web

- On SRX550 devices, the "External storage" option is not supported. Hence, do not select the "External storage" option from the list on the **Maintain > reboot and snapshot** page. [PR741593](#)
- The J-Web interface was vulnerable to HTML cross-site scripting attacks, also called XST or cross-site tracing. [PR752398](#)
- The Layer 2 Transparent Mode feature does not work with group configurations. [PR815225](#)
- In J-Web, if the policy name was "0", the penultimate-hop popping (PHP) function treated it as empty, and traffic log output could not be viewed. [PR853093](#)
- J-Web fails to display the member in the application set after adding it to the nested application set. [PR883391](#)
- On J-Web, when you configure policy, the address set is seen as undefined in the Policy wizard. But, if a policy is created from **Security > Policy > Apply** policy, the address set is seen. [PR892766](#)
- On J-Web, the configured maximum flow memory value key **max-flow-mem** is marked as deprecated and hidden. Therefore, the maximum flow memory value cannot be fetched or displayed in J-Web. [PR894787](#)
- J-Web fails to display all policies under the from or to zone if one of them has the ## string in the description field. [PR917136](#)

License

- On SRX100 High Memory devices, after returning to zero the system licenses are deleted and the device reverts to an SRX100B device. [PR863962](#)

Network Address Translation (NAT)

- NAT-T might not work when the VPN is with Cisco and if the VPN is initiated from a Cisco peer. The VPN negotiates using port UDP 500 instead of UDP 4500 when NAT is involved. [PR869458](#)
- On devices in a chassis cluster, the chassis cluster rule number of sessions in the SNMP query or walk result is the sum of the real number of sessions of the primary node and the secondary node. [PR908206](#)

Security

- The glob implementation in libc allows authenticated remote users to cause a denial of service (CPU and memory consumption) via crafted glob expressions that do not match any pathnames. This vulnerability can be exploited against a device running

Junos OS with FTP services enabled to launch a high CPU utilization partial denial of service attack.

Please refer to JSA10598 for additional information. [PR558494](#)

- If Proxy ARP is enabled on an unnumbered interface, an attacker can poison the ARP cache and create a bogus forwarding table entry for an IP address, effectively creating a denial of service for that subscriber or interface. When Proxy ARP is enabled on an unnumbered interface, the router will answer any ARP message from any IP address which could lead to exploitable information disclosure.

Please refer to JSA10595 for additional information. [PR842092](#)

Switching

- On SRX650 devices, the dot1x:mode:Multiple:Suplicants were authenticated even after a disconnect message was sent from the RADIUS server. [PR786731](#)

Unified Threat Management (UTM)

- The antivirus fallback block notification displays invalid notification option. [PR787063](#)
- When full file-based scanning of antivirus is enabled with Kaspersky scanning, some websites are not accessible. [PR853516](#)
- The flowd process might crash when traffic is processed by UTM. [PR854880](#)
- The device tries to resolve and connect to **cpa.surfcpa.com** and **update.juniper-updates.net** even if there are no licenses or configuration related to UTM. [PR856128](#)
- On all branch SRX Series devices using EWF, a small percentage of the connections to the Websense ThreatSeeker Intelligence Cloud might time out. [PR860514](#)
- The EWF parser mishandled URL and hosts from the HTTP header. This results in an uncategorized EWF reply. [PR862602](#)
- On all branch SRX Series devices with UTM content filtering configured, a long file name encoded with the ISO-2022 might incorrectly match the content filtering extension blocking policy even if the extension blocking list does not contain the type of file extension. As a result, the file is dropped. [PR865607](#)
- On all branch SRX Series devices, new categories for EWF have been added. [PR866160](#)

User Interface and Configuration

- On SRX240 devices (with H2 and B2 model numbers) running Junos OS Release 11.4R8 or 11.4R9, you cannot upgrade to Junos OS Release 11.4R10 or later.

You can upgrade from Junos OS Release 11.4R8 or 11.4R9 to Junos OS Release 12.1X44-D10, 12.1X45-D10, and 12.1X46-D10. [PR934393](#)

Virtual Private Networks (VPNs)

- On an SRX Series device, when a session is closed because the user for that session has signed out from the Junos Pulse, the session close log shows the role information as "N/A". [PR689607](#)
- The SRX Series cluster is used as a VPN concentrator that is connected to remote VPN clients. The Internet key exchange process (daemon) tries to reuse the IP address that was previously assigned to an XAuth client. But the IKEd Xauth attributes are overwritten when the authentication reply is received from Authd. This causes the IKEd to assign a new IP address every time a Phase 1 Security Association (SA) is negotiated. As a result, multiple remote clients cannot connect through VPN. [PR854922](#)
- On all branch SRX Series devices, the Junos Pulse client has been updated from Release 2.0R3 to 4.0R2. [PR868101](#)
- On all branch SRX Series devices, a memory leak occurs on the data plane during continuous interface flapping, such as when interfaces are continuously added or deleted. [PR898731](#)
- For IKEv2, if an SRX Series device running Junos OS Release 12.1X46-D10 is in negotiation with a peer SRX Series device running Junos OS Release 11.4 or 12.1X44, a kmd core file might be generated on the peer device during IPsec child SA rekey. This does not impact any IKEv1 scenarios. [PR915376](#)
- On all branch SRX Series devices configured with group VPN, the flowd process might crash when group VPN Security Association (SA) rekeys and swaps to the new VPN tunnel. [PR925107](#)

Related Documentation

- [New and Changed Features on page 5](#)
- [Changes in Behavior and Syntax on page 17](#)
- [Known Behavior on page 33](#)
- [Known Issues on page 53](#)
- [Documentation Updates on page 76](#)
- [Migration, Upgrade, and Downgrade Instructions on page 86](#)

Documentation Updates

This section lists the errata and changes in Junos OS Release 12.1X46 documentation.

Documentation Updates for the Junos OS Software Documentation

This section lists improvements and outstanding issues with the software documentation.

Junos OS for SRX Series Documentation

The Junos OS for SRX Series technical documentation set has been expanded, restructured, and retitled in Junos OS Release 12.1X46-D10 to make it more comprehensive, easy-to-use, and intuitive. Highlights:

- (New) The Complete Software Guide consolidates all of the release-specific content that applies to Junos OS for SRX Series devices (except release notes) into a three volume set of PDFs that you can download and view offline. The first volume contains getting started and administration information; the second contains feature information; the third contains developer information. You can find the PDFs in the Downloads box on the right side of the *Junos OS for SRX Series Services Gateways, Release 12.1X46* index page.
- (New) The *Getting Started Guide for Branch SRX Series* describes how to get up and running with branch SRX Series devices.
- (Expanded) The *Junos OS Monitoring and Troubleshooting Library for Security Devices* contains significantly more content to help network and security managers keep their SRX Series devices running smoothly in their production environments.
- (Expanded) The *Junos OS for SRX Series Services Gateways, Release 12.1X46* index page has been expanded to serve as a “one stop shop” for all of your Junos OS for SRX Series technical documentation needs.

Junos OS Release Notes

In Junos OS 12.1X46-D10 Release Notes and Junos OS 12.X46-D15 Maintenance Release Notes, the SCCP ALG feature description has the following incorrect information:

Support for SCCP v20—This feature is supported on all SRX Series devices.

Starting in Junos OS Release 12.1X46-D10, the SCCP ALG supports version 20. In SCCP v20, several SCCP messages have been updated with a new format.

The correct information is as follows:

Support for SCCP v20—This feature is supported on all SRX Series devices. Starting in Junos OS Release 12.1X46-D10, the SCCP ALG supports SCCP versions 16, 17, and 20 and several SCCP messages have been updated with a new format. Cisco Call Manager (CM) version 7 uses SCCP version 20.

Administration Guide for Security Devices

- Under the Configuration tab, the “Minimum DHCP Local Server Configuration” topic has been updated to replace the pool name and group name with more appropriate names. The text should read as follows:

```
[edit access]
address-assignment {
  pool acmenetwork family inet {
```

```
        network 192.168.1.0/24;
    }
}

[edit system services]
dhcp-local-server {
    group mobileusers {
        interface ge-1/0/1.0
    }
}

[edit interfaces ge-1/0/1 unit 0]
family {
    inet {
        address 192.168.1.1/24
    }
}
```

BGP Feature Guide for Security Devices

- In “Example: Configuring Route Authentication for BGP,” the following configuration steps in the CLI quick configuration and in the step-by-step procedure sections are not supported on SRX Series devices:

```
set security authentication-key-chains key-chain bgp-auth tolerance 30
set security authentication-key-chains key-chain bgp-auth key 0 secret
this-is-the-secret-password
set security authentication-key-chains key-chain bgp-auth key 0 start-time
2011-6-23.20:19:33-0700
set security authentication-key-chains key-chain bgp-auth key 1 secret
this-is-another-secret-password
set security authentication-key-chains key-chain bgp-auth key 1 start-time
2012-6-23.20:19:33-0700
```

Chassis Cluster Feature Guide for Security Devices

- Under the Configuration tab, in the Example: Configuring an SRX Series Services Gateway for the Branch as a Chassis Cluster, there is a correction in Table 2: SRX Series Services Gateways fxp0 and fxp1 Interfaces Mapping. For the SRX210, the fxp0 Interface should not be ge-0/0/0; it should be fe-0/0/6.
- The **set chassis cluster cluster-id cluster-id node node reboot** operational mode command is missing from the Administration tab. This operational mode command sets the chassis cluster identifier (ID) and node ID on each device, and reboots the devices to enable clustering. This command has two options: **cluster-id cluster-id** (0 through 255) and **node node** (0 or 1). The system uses the chassis cluster ID and chassis cluster node ID to apply the correct configuration for each node (for example, when you use the **apply-groups** command to configure the chassis cluster management interface). The chassis cluster ID and node ID statements are written to the EPROM, and the statements take effect when the system is rebooted. Setting a cluster ID to 0 is equivalent to disabling a cluster. Support for extended cluster identifiers (more than 15 identifiers) added in Junos OS Release 12.1X46-D10. A cluster ID greater than 15 can only be set when the fabric and control link interfaces are connected back-to-back. The command has the following privilege level: maintenance.

If you have a cluster set up and running with an earlier release of Junos OS, you can upgrade to Junos OS Release 12.1X46-D10 or later and re-create a cluster with cluster IDs greater than 16. If for any reason you decide to revert to the previous version of Junos OS that did not support extended cluster IDs, the system comes up with standalone devices after you reboot. If the cluster ID set is less than 16 and you roll back to a previous release, the system comes back with the previous setup.

Interfaces and Routing

- The “Example: Configuring a Serial Interface” of the “Modem Interfaces” guide provides the following incorrect output sample for the **show interfaces se-1/0/0** command:

```
encapsulation ppp;
unit 0 {
  amily inet {
    amily inet;
  }
}
```

The correct output sample is:

```
encapsulation ppp;
unit 0 {
  family inet {
    address 10.10.10.10/24;
  }
}
```

J Series Services Router Advanced WAN Access Configuration Guide

- The example given in the “Configuring Full-Cone NAT” section in the guide available at <http://www.juniper.net/techpubs/software/jseries/junos85/index.html> is incorrect. The correct and updated example is given in the revised guide available at <http://www.juniper.net/techpubs/software/jseries/junos90>).

J2320, J2350, J4350, and J6350 Services Router Getting Started Guide

- The “Connecting to the CLI Locally” section states that the required adapter type is DB-9 female to DB-25 male. This is incorrect; the correct adapter type is DB-9 male to DB-25 male.

J-Web

- **J-Web Security Package Update Help page**—This Help page does not contain information about the download status.
- **J-Web pages for stateless firewall filters**—There is no documentation describing the J-Web pages for stateless firewall filters. To find these pages in J-Web, go to **Configure>Security>Firewall Filters**, and then select **IPv4 Firewall Filters** or **IPv6 Firewall Filters**. After configuring the filters, select **Assign to Interfaces** to assign your configured filters to interfaces.

Junos OS CLI User Guide

- In the **log-prefix** topic, SRX Series is missing from the list of supported platforms and release information.

Modem Interfaces Feature Guide for Security Devices

- The Example: Configuring the 3G Wireless Modem Interface in Modem Interfaces Guide provides the following incorrect information for configuring a dialer filter for the 3G wireless modem interface:
 - **user@host# set firewall family inet dialer-filter corporate-traffic-only term term1 from source-address 20.20.90.4/32**
 - **user@host# set firewall family inet dialer-filter corporate-traffic-only term term1 from destination-address 200.200.201.1/32**
 - **user@host# set firewall family inet dialer-filter corporate-traffic-only term term1 then note**

The following incorrect configuration output is included:

```
[edit]
user@host# show firewall family inet dialer-filter corporate-traffic-only
term term1 {
  from {
    source-address {
      20.20.90.4/32;
    }
    destination-address {
      200.200.201.1/32;
    }
  }
  then note;
}
```

The correct configuration is:

```
user@host# set firewall family inet dialer-filter corporate-traffic-only term term1 then
note
```

The following configuration is output from the correct configuration:

```
[edit]
user@host# show firewall
family inet {
  dialer-filter corporate-traffic-only {
    term term-1 {
      then note;
    }
  }
}
```


Multicast Feature Guide for Security Devices

- Multicast Source Discovery Protocol (MSDP) is not supported on SRX Series devices in any type of custom routing instance.

Network Address Translation

The command **show security nat source persistent-nat-table** under Network Address Translation > Administration > Source NAT operational commands has the following errors:

- The command is missing the **summary** option—Display persistent NAT bindings summary.
- The command contains incomplete sample output—The corrected sample output is as follows:

show security nat source persistent-nat-table internal-ip internal-port

```
user@host> show security nat source persistent-nat-table internal-ip 9.9.9.1 internal-port 60784
```

Internal	Reflective	Source	Type
Left_time/ Curr_Sess_Num/ Source			
In_IP In_Port I_Proto Ref_IP Ref_Port R_Proto NAT Pool			
Conf_time Max_Sess_Num NAT Rule			
9.9.9.1 60784 udp 66.66.66.68 60784	udp	dynamic-customer-source	
any-remote-host 254/300 0/30 105			

show security nat source persistent-nat-table all

```
user@host> show security nat source persistent-nat-table all
```

Internal	Reflective	Source	Type
Left_time/ Curr_Sess_Num/ Source			
In_IP In_Port I_Proto Ref_IP Ref_Port R_Proto NAT Pool			
Conf_time Max_Sess_Num NAT Rule			
9.9.9.1 63893 tcp 66.66.66.68 63893	tcp	dynamic-customer-source	
any-remote-host 192/300 0/30 105			
9.9.9.1 64014 udp 66.66.66.68 64014	udp	dynamic-customer-source	
any-remote-host 244/300 0/30 105			
9.9.9.1 60784 udp 66.66.66.68 60784	udp	dynamic-customer-source	
any-remote-host 254/300 0/30 105			
9.9.9.1 57022 udp 66.66.66.68 57022	udp	dynamic-customer-source	
any-remote-host 264/300 0/30 105			
9.9.9.1 53009 udp 66.66.66.68 53009	udp	dynamic-customer-source	
any-remote-host 268/300 0/30 105			
9.9.9.1 49225 udp 66.66.66.68 49225	udp	dynamic-customer-source	
any-remote-host 272/300 0/30 105			
9.9.9.1 52150 udp 66.66.66.68 52150	udp	dynamic-customer-source	
any-remote-host 274/300 0/30 105			
9.9.9.1 59770 udp 66.66.66.68 59770	udp	dynamic-customer-source	
any-remote-host 278/300 0/30 105			
9.9.9.1 61497 udp 66.66.66.68 61497	udp	dynamic-customer-source	
any-remote-host 282/300 0/30 105			
9.9.9.1 56843 udp 66.66.66.68 56843	udp	dynamic-customer-source	
any-remote-host -/300 1/30 105			

show security nat source persistent-nat-table summary

```

user@host> show security nat source persistent-nat-table summary
Persistent NAT Table Statistics on FPC5 PIC0:
binding total : 65536
binding in use : 0
enode total : 524288
enode in use : 0

```

Routing Protocols Overview for Security Devices

- The default route preference value in the “Understanding Route Preference Values” topic for Static and Static LSPs lists the values incorrectly. The correct values are as follows:

How Route Is Learned	Default Preference
Static	5
Static LSPs	6

Security Policy Applications Feature Guide for Security Devices

- The **show security policies** command output description is missing the definition for the following **Policy statistics** fields:
 - Output packets**—The total number of packets actually processed by the device.
 - Session rate**—The total number of active and deleted sessions.
- On the Overview tab, under IP-Related Predefined Policy Applications, in the topic entitled “Understanding IP-Related Predefined Policy Applications,” the Port column for both TCP-ANY and UDP-ANY should indicate 0-65535. The lead-in sentence should read, “Each entry includes the port and a description of the application.” TCP-ANY means any application that is using TCP, so there is no default port for it. The same is true for UDP-ANY.
- In the topic entitled “Understanding Miscellaneous Predefined Policy Applications,” table “Predefined Miscellaneous Applications” is incomplete. Under the RADIUS row, add a new row:

Table 6: Predefined Miscellaneous Applications

Application	Port	Description
RADIUS Accounting	1813	Enables the collecting of statistical data about users logging in to or out from a LAN and sending the data to a RADIUS Accounting server.

In table “Predefined Miscellaneous Applications” replace the IPsec-NAT row with the following:

Table 7: Predefined Miscellaneous Applications

Application	Port	Description
IKE	500	Internet Key Exchange is the protocol that sets up a security association in the IPsec protocol suite.
IKE-NAT	4500	Helps to perform Layer 3 NAT for S2C IKE traffic.

Application	Port	Description
VoIP	389	Internet Locator Service (ILS)
	522	User Location Service (ULS)
	1503	T.120 Data sharing
	1719	H.225 RAS message
	1720	Q.931 Call Setup
	1731	Audio Call Control
	5060	SIP protocol

Various Guides

- Some Junos OS user, reference, and configuration guides—for example the [Junos Software Routing Protocols Configuration Guide](#), [Junos OS CLI User Guide](#), and [Junos OS System Basics Configuration Guide](#)—mistakenly do not indicate SRX Series device support in the “Supported Platforms” list and other related support information; however, many of those documented Junos OS features are supported on SRX Series devices. For full, confirmed support information about SRX Series devices, please refer to Feature Explorer:
<http://pathfinder.juniper.net/feature-explorer/select-software.html?swName=Junos+OS&typ=1>.

WLAN Feature Guide for Security Devices

- This guide is missing information that the AX411 Access Point can be managed from SRX100 and SRX110 devices.
- This guide is missing the information that on all branch SRX Series devices, managing AX411 WLAN Access Points through a Layer 3 Aggregated Ethernet (ae) interface is not supported.

Documentation Updates for the Junos OS Hardware Documentation

This section lists outstanding issues with the hardware documentation.

J Series Services Routers Hardware Guide

- The procedure “Installing a DRAM Module” omits the following condition:
All DRAM modules installed in the router must be the same size (in megabytes), type, and manufacturer. The router might not work properly when DRAM modules of different sizes, types, or manufacturer are installed.
- This guide incorrectly states that only the J2350 Services Router complies with NEBS criteria. It should state that the J2350, J4350, and J6350 routers comply with NEBS criteria.
- This guide is missing information about 100Base-LX connector support for 1-port and 6-port Gigabit Ethernet uPIMs.

SRX Series Services Gateways for the Branch Physical Interface Modules Hardware Guide

- This guide incorrectly documents that slot 3 of the SRX550 Services Gateway can be used to install GPIMs. The correct information is:
 - In Table 10: “SRX Series Services Gateway Interface Port Number Examples”, for 2-Port 10 Gigabit Ethernet XPIM, you can install the XPIM only in slot 6 of the SRX550 Services Gateway.
 - In Table 44: “Slots for 20-Gigabit GPIMs, for 20-Gigabit GPIM slots”, you can install the GPIM only in slot 6 of the SRX550 Services Gateway.

SRX100 Services Gateway Hardware Guide

- In the “Connecting an SRX100 Services Gateway to the J-Web Interface” section, the following information is missing in the note:



NOTE: Microsoft Internet Explorer version 6.0 is also supported as backward compatible from Microsoft Internet Explorer version 7.0.

SRX210 Services Gateway Hardware Guide

- In the “Connecting an SRX210 Services Gateway to the J-Web Interface” section, the following information is missing in the note:



NOTE: Microsoft Internet Explorer version 6.0 is also supported as backward compatible from Microsoft Internet Explorer version 7.0.

- The “SRX210 Services Gateway Specifications” table lists the values for chassis height, chassis width, chassis depth, chassis weight, and noise level incorrectly. The correct values are as follows:
 - Chassis height—1.73 in. (44 mm)
 - Chassis width—11.02 in. (280 mm)

- Chassis depth—7.13 in. (181 mm)
- Chassis weight:
 - 3.46 lb (1.57 kg) for SRX210 Services Gateway without PoE (no interface modules)
 - 3.55 lb (1.61 kg) for SRX210 Services Gateway with PoE (no interface modules)
- Noise level—29.1 dB per EN ISO 7779

SRX220 Services Gateway Hardware Guide

- The “SRX220 Services Gateway Specifications” table lists the values for chassis height, chassis width, chassis depth, chassis weight, and noise level incorrectly. The correct values are as follows:
 - Chassis height—1.73 in. (44 mm)
 - Chassis width—14.29 in. (363 mm)
 - Chassis depth—7.13 in. (181 mm)
 - Chassis weight:
 - 4.52 lb (2.05 kg) for SRX220 models without PoE (no interface modules)
 - 4.62 lb (2.10 kg) for SRX220 models with PoE (no interface modules)
 - Noise level—51.1 dB per EN ISO 7779

SRX240 Services Gateway Hardware Guide

- In the “Connecting the SRX240 Services Gateway to the J-Web Interface” section, the following information is missing in the note:



NOTE: Microsoft Internet Explorer version 6.0 is also supported as backward compatible from Microsoft Internet Explorer version 7.0.

- The “Maintaining the SRX650 Services Gateway Power Supply” section incorrectly states that the status of the power supplies on the SRX650 Services Gateway can be checked by issuing the **show chassis environment pem** command. The **show chassis environment pem** command is not supported on the SRX650 Services Gateway.

SRX110 Services Gateway 3G USB Modem Quick Start

- The SRX110 Services Gateway 3G USB Modem Quick Start has been updated with the J-Web procedures, and it is available on the Juniper Networks website.

SRX210 Services Gateway 3G ExpressCard Quick Start

- Several tasks are listed in the wrong order. “Task 6: Connect the External Antenna” should appear before “Task 3: Check the 3G ExpressCard Status,” because the user needs to connect the antenna before checking the status of the 3G ExpressCard. The correct order of the tasks is as follows:

1. Install the 3G ExpressCard
 2. Connect the External Antenna
 3. Check the 3G ExpressCard Status
 4. Configure the 3G ExpressCard
 5. Activate the 3G ExpressCard Options
- In “Task 6: Connect the External Antenna,” the following sentence is incorrect and redundant: “The antenna has a magnetic mount, so it must be placed far away from radio frequency noise sources including network components.”
 - In the “Frequently Asked Questions” section, the answer to the following question contains an inaccurate and redundant statement:

Q: Is an antenna required? How much does it cost?

A: The required antenna is packaged with the ExpressCard in the SRX210 Services Gateway 3G ExpressCard kit at no additional charge. The antenna will have a magnetic mount with ceiling and wall mount kits within the package.

In the answer, the sentence “The antenna will have a magnetic mount with ceiling and wall mount kits within the package” is incorrect and redundant.

SRX210 Services Gateway Quick Start Guide

- The section on installing software packages is missing the following information:

On SRX210 devices, the `/var` hierarchy is hosted in a separate partition (instead of the `root` partition). If Junos OS installation fails as a result of insufficient space:
 1. Use the **request system storage cleanup** command to delete temporary files.
 2. Delete any user-created files both in the `root` partition and under the `/var` hierarchy.

Related Documentation

- [New and Changed Features on page 5](#)
- [Changes in Behavior and Syntax on page 17](#)
- [Known Behavior on page 33](#)
- [Known Issues on page 53](#)
- [Resolved Issues on page 54](#)
- [Migration, Upgrade, and Downgrade Instructions on page 86](#)

Migration, Upgrade, and Downgrade Instructions

This section includes the following topics:

- [Upgrading and Downgrading among Junos OS Releases on page 87](#)
- [Upgrading an AppSecure Device on page 89](#)
- [Network and Security Manager Support on page 89](#)

- [Upgrade and Downgrade Scripts for Address Book Configuration on page 89](#)
- [Hardware Requirements on page 92](#)

Upgrading and Downgrading among Junos OS Releases

All Junos OS releases are listed in sequence on the JUNOS Software Dates & Milestones webpage:

<http://www.juniper.net/support/eol/junos.html>

To help in understanding the examples that are presented in this section, a portion of that table is replicated here. See [Table 8 on page 87](#). Note that releases footnoted with a 1 are Extended End-of-Life (EEOL) releases.

Table 8: Junos Software Dates & Milestones

Product	FRS Date
Junos 12.1X44 ¹²	01/18/2013
Junos 12.1	03/28/2012
Junos 11.4 ¹	12/21/2011
Junos 11.3	08/15/2011
Junos 11.2	08/03/2011
Junos 11.1	03/29/2011
Junos 10.4 ¹	12/08/2010
Junos 10.3	08/15/2010
Junos 10.2	05/28/2010
Junos 10.1	02/15/2010
Junos 10.0 ¹	11/04/2009
Junos 9.6	08/06/2009
Junos 9.5	04/14/2009
Junos 9.4	02/11/2009
Junos 9.3 ¹	11/14/2008
Junos 9.2	08/12/2008
Junos 9.1	04/28/2008

Table 8: Junos Software Dates & Milestones *(continued)*

Product	FRS Date
Junos 9.0	02/15/2008
Junos 8.5 ¹	fwd-srns-context

You can directly upgrade or downgrade between any two Junos OS releases that are within three releases of each other.

- Example: Direct release upgrade

Release 10.3 → (*bypassing Releases 10.4 and 11.1*) Release 11.2

To upgrade or downgrade between Junos OS releases that are more than three releases apart, you can upgrade or downgrade first to an intermediate release that is within three releases of the desired release, and then upgrade or downgrade from that release to the desired release.

- Example: Multistep release downgrade

Release 11.3 → (*bypassing Releases 11.2 and 11.1*) Release 10.4 → Release 10.3

Juniper Networks has also provided an even more efficient method of upgrading and downgrading using the Junos OS EEOL releases. EEOL releases generally occur once a calendar year and can be more than three releases apart. For a list of, EEOL releases, go to <http://www.juniper.net/support/eol/junos.html>

You can directly upgrade or downgrade between any two Junos OS EEOL releases that are within three EEOL releases of each other.

- Example: Direct EEOL release upgrade

Release 9.3 (EEOL) → (*bypassing Releases 10.0 [EEOL] and 10.4 [EEOL]*) Release 11.4 (EEOL)

To upgrade or downgrade between Junos OS EEOL releases that are more than three EEOL releases apart, you can upgrade first to an intermediate EEOL release that is within three EEOL releases of the desired EEOL release, and then upgrade from that EEOL release to the desired EEOL release.

- Example: Multistep release upgrade using intermediate EEOL release

Release 8.5 (EEOL) → (*bypassing Releases 9.3 [EEOL] and 10.0 [EEOL]*) Release 10.4 (EEOL) → Release 11.4 (EEOL)

You can even use a Junos OS EEOL release as an intermediate upgrade or downgrade step if your desired release is several releases later than your current release.

- Example: Multistep release upgrade using intermediate EEOL release

Release 9.6 → Release 10.0 (EEOL) → Release 10.2

For additional information about how to upgrade and downgrade, see the *Junos OS Installation and Upgrade Guide*.

Upgrading an AppSecure Device

Use the no-validate Option for AppSecure Devices.

For devices implementing AppSecure services, use the no-validate option when upgrading from Junos OS Release 11.2 or earlier to Junos OS 11.4R1 or later. The application signature package used with AppSecure services in previous releases has been moved from the configuration file to a signature database. This change in location can trigger an error during the validation step and interrupt the Junos OS upgrade. The no-validate option bypasses this step.

Network and Security Manager Support

Network and Security Manager (NSM) support for SRX Series Services Gateways and J Series Services Routers with Junos OS 12.1X46-D10 is available only with NSM versions 2012.2R6 / 2012.1R10 and later. For additional information, see [Network and Security Manager](#) documentation.

Upgrade and Downgrade Scripts for Address Book Configuration

Beginning with Junos OS Release 12.1, you can configure address books under the **[security]** hierarchy and attach security zones to them (zone-attached configuration). In Junos OS Release 11.1 and earlier, address books were defined under the **[security zones]** hierarchy (zone-defined configuration).

You can either define all address books under the **[security]** hierarchy in a zone-attached configuration format or under the **[security zones]** hierarchy in a zone-defined configuration format; the CLI displays an error and fails to commit the configuration if you configure both configuration formats on one system.

Juniper Networks provides Junos operation scripts that allow you to work in either of the address book configuration formats (see [Figure 1 on page 91](#)).

- [About Upgrade and Downgrade Scripts on page 89](#)
- [Running Upgrade and Downgrade Scripts on page 91](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases on page 92](#)

About Upgrade and Downgrade Scripts

After downloading Junos OS Release 12.1, you have the following options for configuring the address book feature:

- **Use the default address book configuration**—You can configure address books using the zone-defined configuration format, which is available by default. For information on how to configure zone-defined address books, see the Junos OS Release 11.1 documentation.

- **Use the upgrade script**—You can run the upgrade script available on the Juniper Networks support site to configure address books using the new zone-attached configuration format. When upgrading, the system uses the zone names to create address books. For example, addresses in the trust zone are created in an address book named **trust-address-book** and are attached to the trust zone. IP prefixes used in NAT rules remain unaffected.

After upgrading to the zone-attached address book configuration:

- You cannot configure address books using the zone-defined address book configuration format; the CLI displays an error and fails to commit.
- You cannot configure address books using the J-Web interface.


For information on how to configure zone-attached address books, see the Junos OS Release 12.1 documentation.

- **Use the downgrade script**—After upgrading to the zone-attached configuration, if you want to revert to the zone-defined configuration, use the downgrade script available on the Juniper Networks support site. For information on how to configure zone-defined address books, see the Junos OS Release 11.1 documentation.



NOTE: Before running the downgrade script, make sure to revert any configuration that uses addresses from the global address book.

Figure 1: Upgrade and Downgrade Scripts for Address Books



Running Upgrade and Downgrade Scripts

The following restrictions apply to the address book upgrade and downgrade scripts:

- The scripts cannot run unless the configuration on your system has been committed. Thus, if the zone-defined address book and zone-attached address book configurations are present on your system at the same time, the scripts will not run.
- The scripts cannot run when the global address book exists on your system.
- If you upgrade your device to Junos OS Release 12.1 and configure logical systems, the master logical system retains any previously configured zone-defined address book configuration. The master administrator can run the address book upgrade script to convert the existing zone-defined configuration to the zone-attached configuration. The upgrade script converts all zone-defined configurations in the master logical system and user logical systems.



NOTE: You cannot run the downgrade script on logical systems.

For information about implementing and executing Junos operation scripts, see the *Junos OS Configuration and Operations Automation Guide*.

Upgrade and Downgrade Support Policy for Junos OS Releases

Support for upgrades and downgrades that span more than three Junos OS releases at a time is not provided, except for releases that are designated as Extended End-of-Life (EEOL) releases. EEOL releases provide direct upgrade and downgrade paths—you can upgrade directly from one EEOL release to the next EEOL release even though EEOL releases generally occur in increments beyond three releases.

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 10.0, 10.4, and 11.4 are EEOL releases. You can upgrade from Junos OS Release 10.0 to Release 10.4 or even from Junos OS Release 10.0 to Release 11.4. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind. For example, you cannot directly upgrade from Junos OS Release 10.3 (a non-EEOL release) to Junos OS Release 11.4 or directly downgrade from Junos OS Release 11.4 to Junos OS Release 10.3.

To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

For more information about EEOL releases and to review a list of EEOL releases, see <http://www.juniper.net/support/eol/junos.html>.

Hardware Requirements

Transceiver Compatibility for SRX Series and J Series Devices

We strongly recommend that only transceivers provided by Juniper Networks be used on SRX Series and J Series interface modules. Different transceiver types (long-range, short-range, copper, and others) can be used together on multiport SFP interface modules as long as they are provided by Juniper Networks. We cannot guarantee that the interface module will operate correctly if third-party transceivers are used.

Please contact Juniper Networks for the correct transceiver part number for your device.

Power and Heat Dissipation Requirements for J Series PIMs

On J Series Services Routers, the system monitors the PIMs and verifies that the PIMs fall within the power and heat dissipation capacity of the chassis. If power management is enabled and the capacity is exceeded, the system prevents one or more of the PIMs from becoming active.



CAUTION: Disabling the power management can result in hardware damage if you overload the chassis capacities.

You can also use CLI commands to choose which PIMs are disabled. For details about calculating the power and heat dissipation capacity of each PIM and for troubleshooting procedures, see the *J Series Services Routers Hardware Guide*.

Supported Third-Party Hardware

The following third-party hardware is supported for use with J Series Services Routers running Junos OS.

- **USB Modem**

We recommend using a U.S. Robotics USB 56K V.92 Modem, model number USR 5637.

- **Storage Devices**

The USB slots on J Series Services Routers accept a USB storage device or USB storage device adapter with a CompactFlash card installed, as defined in the *CompactFlash Specification* published by the CompactFlash Association. When the USB device is installed and configured, it automatically acts as a secondary boot device if the primary CompactFlash card fails on startup. Depending on the size of the USB storage device, you can also configure it to receive any core files generated during a router failure. The USB device must have a storage capacity of at least 256 MB.

[Table 9 on page 93](#) lists the USB and CompactFlash card devices supported for use with the J Series Services Routers.

Table 9: Supported Storage Devices on the J Series Services Routers

Manufacturer	Storage Capacity	Third-Party Part Number
SanDisk—Cruzer Mini 2.0	256 MB	SDCZ2-256-A10
SanDisk	512 MB	SDCZ3-512-A10
SanDisk	1024 MB	SDCZ7-1024-A10
Kingston	512 MB	DTI/512KR
Kingston	1024 MB	DTI/1GBKR
SanDisk—ImageMate USB 2.0 Reader/Writer for CompactFlash Type I and II	N/A	SDDR-91-A15
SanDisk CompactFlash	512 MB	SDCFB-512-455
SanDisk CompactFlash	1 GB	SDCFB-1000.A10

J Series CompactFlash and Memory Requirements

[Table 10 on page 94](#) lists the CompactFlash card and DRAM requirements for J Series Services Routers.

Table 10: J Series CompactFlash Card and DRAM Requirements

Model	Minimum CompactFlash Card Required	Minimum DRAM Required	Maximum DRAM Supported
J2320	1 GB	1 GB	1 GB
J2350	1 GB	1 GB	1 GB
J4350	1 GB	1 GB	2 GB
J6350	1 GB	1 GB	2 GB

Related Documentation

- [New and Changed Features on page 5](#)
- [Changes in Behavior and Syntax on page 17](#)
- [Known Behavior on page 33](#)
- [Known Issues on page 53](#)
- [Resolved Issues on page 54](#)
- [Documentation Updates on page 76](#)

Junos OS Release Notes for High-End SRX Series

Powered by Junos OS, Juniper Networks high-end SRX Series Services Gateways provide robust networking and security services. High-end SRX Series Services Gateways are designed to secure enterprise infrastructure, data centers, and server farms. The high-end SRX Series Services Gateways include the SRX1400, SRX3400, SRX3600, SRX5400, SRX5600, and SRX5800 devices.

- [New and Changed Features on page 95](#)
- [Changes in Behavior and Syntax on page 113](#)
- [Known Behavior on page 135](#)
- [Known Issues on page 155](#)
- [Resolved Issues on page 156](#)
- [Documentation Updates on page 183](#)
- [Migration, Upgrade, and Downgrade Instructions on page 191](#)

New and Changed Features

The following features have been added to Junos OS Release 12.1X46. Following the description is the title of the manual or manuals to consult for further information.



.....
NOTE: For the latest updates about support and issues on Junos Pulse, see the Release 12.1X46-D15 Software Features
.....

Release 12.1X46-D30 Software Features

Application Layer Gateways (ALGs)

- **MS-RPC ALG and Sun RPC ALG map table scaling for SRX Series devices**— Starting with Junos OS Release 12.1x46-D30, the MS-RPC ALG and Sun RPC ALG dynamically allocate new mapping entries instead of using a default size (512 entries). They also offer a flexible time-based RPC mapping entry that removes the mapping entry (auto-clean) without affecting the associated active RPC sessions, including both control session and data session.

[See [RPC ALG Feature Guide for Security Devices](#).]

Release 12.1X46-D25 Software Features

General Packet Radio Service Feature Guide for Security Devices

- **GTP GSN Table Ager [High-end SRX Series]**—One SRX Series device supports a total of 36,000 GSN entries, each of which was saved permanently prior to this release. To prevent GSN entry exhaustion caused by frequent short-time roaming among countries, visiting GSNs are recorded when subscribers access the home GPRS core network from visiting countries. These entries are not deleted when the subscribers return home, but no further traffic is passed. The GTP GSN table ager causes the idling GSN entries to time out, preventing inactive GSNs from taking up too much space.

[See the “General Packet Radio Service” section in the Junos OS 12.1X46-D25 Feature Guide.]

TCP/TLS support for real time logging

- **TCP/TLS support for real-time logging [High-end SRX Series]**—Starting in Junos OS Release 12.1X46-D25, a secure mechanism, enabled through a plug-in during system initialization, encrypts and transports dataplane syslog messages to TLS-capable syslog receivers (such as the Juniper Networks STRM or a standards-based third-party device) over TCP. The SPU generates the log data. By default, port 514 is used for TCP logging and port 6514 is used for TLS logging. As a log client, a TCP/TLS connection is initiated to the log server.

[See the “Syslog Messages” section in the Junos OS 12.1X46-D25 Release Feature Guide.]

Release 12.1X46-D20 Software Features

Chassis Cluster

- **Autorecovery of fabric link [SRX Series]**—The fabric link feature supports autorecovery, which includes the following enhancements:
 - Fabric monitoring feature is enabled by default on high-end SRX Series, and hence recovery of fabric link and synchronization takes place automatically.
 - If the fabric link goes down, RG1+ becomes ineligible on either the secondary node or the node with failures, by default. The node remains in this state until the fabric link comes up or the other node goes away.

- If the fabric link goes down followed by the control link, then after approximately 66 seconds the secondary node (or the node with failures) assumes that the remote node is dead and takes over as the primary node.

[See [Understanding Chassis Cluster Fabric Links](#).]

- **Enhanced debugging support for chassis cluster** [SRX Series]—The chassis cluster debugging functionality has the following enhancements:
 - The **show chassis cluster status** command output includes failure reasons (acronyms and their expansions) when the redundancy group's priority is zero.
 - Cleaner jsrpd process includes removing unwanted logs and moving the debug log message from level LOG_INFO to LOG_DEBUG.
 - The **show chassis cluster information** command output displays redundancy group, LED, and monitored failure details.
 - SNMP traps send messages when a node's weight goes down and also when it recovers.
 - The **show chassis cluster ip-monitoring** command output displays both the global threshold and the current threshold of each node and displays the weight of each monitored IP address.
 - A syslog message appears when the control link goes down.

[See [show chassis cluster ip-monitoring status](#).]

Public Key Infrastructure (PKI)

- **Online Certificate Status Protocol (OCSP)** [SRX Series]—OCSP, like CRL, checks the revocation status of X509 certificates. Requests are sent to the OCSP server(s) configured in a CA profile with the `ocsp url` statement at the `[edit security pki ca-profile profile-name revocation-check]` hierarchy level. The `use-ocsp` option must also be configured. If there is no response from the OCSP server, the request is then sent to the location specified in the certificate's AuthorityInfoAccess extension.

[See the “Public Key Infrastructure (PKI)” section in the [Junos OS 12.1X46-D20 Feature Guide](#).]

Routing Protocols

- **OSPFv3 IPsec authentication and confidentiality** [SRX Series]—OSPF for IPv6, also known as OSPF version 3 (OSPFv3), does not have built-in authentication to ensure that routing packets are not altered and re-sent to the router. In Junos OS Release 12.1X46-D20, IPsec can be used to secure OSPFv3 interfaces and virtual links and provide encryption for OSPF packets.

To configure IPsec for OSPF/OSPFv3, define a security association (SA) with the `security-association sa-name` configuration option at the `[edit security ipsec]` hierarchy level. The configured SA is then applied to the OSPF/OSPFv3 interface or virtual link configuration.

[See the “Routing Protocols” section in the [Junos OS 12.1X46-D20 Feature Guide](#).]

Unified Threat Management (UTM)

- **UTM on next-generation SPC [SRX5400, SRX5600 and SRX5800]**—This feature provides support for UTM features, including Sophos antivirus, content filtering, antispam, and enhanced Web filtering on next-generation SPCs.
- **UTM license enforcement [SRX Series]**—License enforcement is supported for UTM features, including Sophos antivirus, enhanced Web filtering, and antispam filtering on all high-end SRX Series devices in addition to branch SRX Series devices. You can add or remove UTM licenses on SRX Series devices. Each feature license is tied to exactly one software feature and is valid for exactly one device.

[Table 1 on page 7](#) lists the license modules and the license names.

Table 11: UTM License Information

UTM Module	License Name
SAV	av_key_sophos_engine
AS	anti_spam_key_sbl
EWf	wf_key_websense_ewf

On branch SRX Series devices, after you install the license and reboot the device, the device reserves more memory for UTM features, and hence decreases the session

capacity. Use the **set security forwarding-process application-services enable-utm-memory** command to manually reallocate the memory for UTM features. You must reboot the device for the configuration to take effect.

[See the “UTM” section in the [Junos OS 12.1X46-D20 Feature Guide](#).]

[See [License Enforcement](#).]

VPNs

- **HMAC-SHA-256-128 authentication** [High-end SRX Series]—Starting with Junos OS Release 12.1X46-D20, HMAC-SHA-256-128 authentication is supported for IPsec proposals and manual security associations on high-end SRX Series devices. You can specify the **hmac-sha-256-128** option at the **[edit security ipsec proposal *proposal-name*]** and the **[edit security ipsec vpn *vpn-name* manual]** hierarchy levels.

[See the “VPNs” section in the [Junos OS 12.1X46-D20 Feature Guide](#).]

Release 12.1X46-D15 Software Features

Routing Protocols

- **OSPFv2 support** [High-end SRX Series]—OSPFv2 interfaces are supported on nonbroadcast multiaccess (NBMA) networks and point-to-point access networks on high-end SRX Series devices.

When you configure OSPFv2 on an NBMA network, 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.

An NBMA interface behaves similarly to a point-to-multipoint interface but requires election and operation of a designated router and a backup designated router.

Use the following CLI commands to configure an OSPFv2 interface on an NBMA or a point-to-multipoint network:

- **set protocols ospf area *area-number* interface *interface-name* neighbor *address-of-neighbor***
- **set protocols ospf area *area-number* interface *interface-name* interface-type *interface-type* (*nbma* or *p2mp*)**

[See “Routing Protocols” section in [Junos OS 12.1X46-D15 Feature Guide](#).]

Release 12.1X46-D10 Software Features

Application Layer Gateways (ALGs)

- **ALG message buffer optimization**—Starting in Junos OS Release 12.1X46-D10, the ALG message buffer optimization feature is enhanced to reduce high memory consumption. This feature is supported on all SRX Series and J Series devices.

A message buffer is allocated only when the packet is ready to process. The buffer is freed after the packet completes ALG handling, including modifying the payload, performing NAT, opening a pinhole for a new connection between a client and a server,

and transferring data between a client and a server located on opposite sides of a Juniper Networks device.

This feature has the following enhancements:

- Unnecessary objcache buffering is avoided, resulting in low memory utilization.
- jbuf manipulation is used to simplify the message buffer logic.
- Full-fledged message buffer support for ALG line breaker is more flexible.
- ALG Manager and ALG plug-in logic clarity are optimized.

[See [alg-manager](#).]

- **IPv6 support for PPTP ALG**—Starting with Junos OS Release 12.X46, this feature is supported on all SRX Series devices.

PPTP ALG provides an ALG for the Point-to-Point Tunneling Protocol (PPTP). The PPTP is a Layer 2 protocol that tunnels PPP data across TCP/IP networks. The PPTP client is freely available on Windows systems and popularly applied on Linux systems; it is widely deployed for building VPNs.

To support IPv6, the PPTP ALG parses both IPv4 and IPv6 PPTP packets, performs NAT, and then opens a pinhole for the data tunnel. The flow module supports IPv6 to parse the GRE packet and use the GRE call ID as fake port information to search the session table and gate table.

[See [PPTP ALG Feature Guide for Security Devices](#).]

- **IPv6 support for RTSP ALG**—This feature is supported on all SRX Series and J Series devices.

RTSP (Real-Time Streaming Protocol) is an Application Layer protocol for controlling the delivery of data with real-time properties. The RTSP ALG accesses existing media files over the network and controls the replay of the media.

Starting with Junos OS Release 12.1X46-D10, IPv6 is supported on the RTSP ALG along with NAT-PT mode and NAT64 address translation.

This feature enables the RTSP ALG to parse IPv6 RTSP packets, open an IPv6 pattern pinhole, and translate the Layer 7 IPv6 address according to the NAT configuration. Also, support for IPv6 RTSP transaction pass through under permission policy and IPv6 RTSP transaction pass through under NAT-PT and NAT 64 are enabled.

[See [SIP RTSP ALG Feature Guide for Security Devices](#).]

- **IPv6 support for SIP ALG**—This feature is supported on all SRX Series and J Series devices.

Starting with Junos OS Release 12.1X46-D10, IPv6 is supported on the SIP ALG along with NAT-PT mode and NAT64 address translation.

The SIP ALG processes the IPv6 address in the same way it processes the IPv4 address for updating the payload if NAT is configured and opening pinholes for future traffic.

NAT-PT is implemented by normal NAT from IPv6 address to IPv4 address and vice versa. The SIP ALG processes those address translation in payload just as the addresses are processed in normal NAT.

NAT64 is a mechanism to allow IPv6 hosts to communicate with IPv4 servers. NAT64 is required to keep the IPv6 to IPv4 address mapping.

Previously Session Traversal Utilities for NAT (STUN) worked without the SIP ALG. This means that the SIP ALG was not involved when persistent NAT was configured.

Starting with Junos OS Release 12.1X46-D10, STUN can coexist with the SIP ALG and SIP ALG is involved when persistent NAT is configured.

[See [SIP ALG Feature Guide for Security Devices](#).]

Chassis Cluster

- **Chassis cluster**—Starting in Junos OS Release 12.1X46-D10, the SRX5K-MPC adds the support of using 40-Gigabit Ethernet and 100-Gigabit Ethernet ports as chassis cluster fabric ports. This feature is supported on the SRX5400, SRX5600, and SRX5800. This enhancement saves one more slot on chassis and also improves chassis cluster fabric link performance. In addition, you can also use 10G port on SRX5K-MPC as fabric port with a 10x10GE MIC installed on it.

[See [Understanding Chassis Cluster Fabric Links](#).]

Dynamic Host Configuration Protocol (DHCP)

- **DHCP relay**—Starting in Junos OS Release 12.1X46-D10, this feature is supported on all high-end SRX Series devices.

The existing DHCP relay feature has been enhanced to include support for high-end SRX Series devices along with chassis cluster support.

You can configure DHCP relay options on the device and enable the device to function as a DHCP relay agent. A DHCP relay agent forwards DHCP request and reply packets between a DHCP client and a DHCP server.

To configure the DHCP relay agent on the device, include the **dhcp-relay** statement at the **[edit forwarding-options]** hierarchy level.

[See [Understanding DHCP Relay Agent Operation](#).]

Flow and Processing

- **Enhanced IPv6 support for the screen feature**—This feature is supported on all high-end SRX Series devices.

IPv6 support is extended for the following screen features:

- IPv6 extension header checking and filtering
- IPv6 packet header checking and filtering
- ICMPv6 checking and filtering

New statements and commands allow you to configure these enhancements using security zones similar to previous screen configurations. You can enable, disable, and update screens to drop packets, create logs, and provide increased statistics for IPv6 traffic.



NOTE: By default, IPv6 packets bypass the screen feature.

[See [Understanding IPv6 Support for Screens](#).]

- **Enhancements to flow trace options**—This feature is supported on all high-end SRX Series devices.

Starting in Junos OS Release 12.1X46-D10, flow trace granularity has been enhanced to filter logs effectively. As a result you can access relevant trace messages easily and avoid large traces that slow down your system. You can set the level of message you want displayed by using the new **trace-level** statement at the **[edit security flow traceoptions]** hierarchy level. And, use new flags to trace additional operations such as fragmentation, high availability, multicast, session, tunnel, and route.

[See [traceoptions \(Security Flow\)](#).]

- **Monitoring flow sessions**—This feature is supported on all high-end SRX Series devices.

Beginning with Junos OS Release 12.1X46-D10, you can monitor flow using filters that match different criteria (such as source and destination addresses). New operational mode commands **monitor security flow filter** and **monitor security flow file** have been added. These commands allow you to debug without having to commit or modify your running configuration. Previously, you were required to commit the configuration to turn on trace options, which could possibly change the state of your device.

[See [Monitoring Security Flow Sessions Overview](#).]

General Packet Radio Service (GPRS)

- **NAT for GPRS tunneling protocol (GTP)**:—Starting in Junos OS Release 12.1X46-D10, static NAT for GTP packets is supported on all high-end SRX Series devices.

This feature has the following enhancements:

- For GTP, control (GTP-C), as part of the GPRS IP address negotiation, embedded IP addresses are included in the packet data protocol (PDP) context request or response messages.
- For GTP, user plane (GTP-U), GTP-U carries encapsulated user payload in an IP packet. When NAT is enabled, only the outer IP packet needs to be translated, the embedded IP addresses will not be translated.

[See [Understanding GTP-U Inspection](#).]

- **GTP unified in-service software upgrade support (ISSU)**—Junos OS Release 12.1X46-D10 adds support for unified ISSU on the GPRS tunneling protocol (GTP). This feature is supported on all high-end SRX Series devices.

GTP supports unified ISSU between two SRX Series devices running two different Junos OS releases. Unified ISSU is applied on a chassis cluster, enabling a software upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic.

[See [Understanding GTP-U Inspection](#).]

Interfaces and Routing

- **Link aggregation**—Starting in Junos OS Release 12.1X46-D10, the SRX5K-MPC supports the LAG and LACP features on the SRX5400, SRX5600, and SRX5800.

The following LAG and LACP features are supported:

- Increases bandwidth, provides graceful degradation as failure occurs, and increases availability.
- Provides network redundancy by load-balancing traffic across all available links. If one of the links should fail, the system automatically load-balances traffic across all remaining links.
- Enables automatic addition and deletion of individual links to the aggregate bundle without user intervention.
- Provides link monitoring to check whether both ends of the bundle are connected to the correct group, enables or disables link protection, configures the LACP interval, and supports centralized and distributed modes.

[See [LAG and LACP Support on the SRX5000 Module Port Concentrator](#).]

[See [Understanding Aggregated Ethernet Interfaces](#).]

IP Spoofing

- **IP spoofing in transparent mode**—Starting in Junos OS Release 12.1X46-D10, this feature is supported on all high-end SRX Series devices.

The IP spoofing feature has been enhanced to include Layer 2 transparent mode support. IP spoofing is most frequently used in denial-of-service attacks. In an IP spoofing attack, the attacker gains access to a restricted area of the network and inserts a false source address in the packet header to make the packet appear to come from a trusted source. When SRX Series devices are operating in transparent mode, the IP spoof-checking mechanism makes use of address book entries.



NOTE: IP spoofing in Layer 2 transparent mode does not support DNS and wildcard addresses.

[See [Understanding IP Spoofing in Layer 2 Transparent Mode](#).]

J-Web

- **Management support for NAT options**—Starting in Junos OS Release 12.1X46-D10, support is provided to monitor the following NAT options on all SRX Series devices:
 - Utilization for all source pools
 - Successful, failed, and current sessions for source pools, source rules, destination rules, and static rules

- Source addresses and source ports for static rules
- Source ports for source rules
- Support is provided to configure the following NAT options on all SRX Series devices:
 - Source address and port as match criteria for static rules
 - Source port as match criteria for source rules
 - Upper and lower thresholds at which an SNMP trap is triggered for source rules and pools, destination rules, and static rules
- **User Firewall J-Web support**
 - **Source identity-based firewall policy**—Starting in Junos OS Release 12.1X46-D10, this feature is supported on the existing Firewall Policies Configuration and Monitoring Policies pages on all high-end SRX Series devices. This feature allows you to configure and monitor source identities in a firewall policy.
 - **Configure firewall authentication integrated with user firewall**—Starting in Junos OS Release 12.1X46-D10, this feature is supported on all high-end SRX Series devices. You use this feature by specifying the access profile and SSL termination profile.
 - **New J-Web pages for user firewall**—Starting in Junos OS Release 12.1X46-D10, new user firewall pages are supported on all high-end SRX Series devices.

The following webpages have been added to the J-Web user interface:

- **Authentication Priority Configuration Page**—You can either disable an optional authentication source or reassign a unique priority to it.
- **Local Authentication Configuration Page and Local Authentication Monitoring Page**—You can configure and monitor local Firewall authentication.
- **UAC Settings Configuration Page and UAC Authentication Monitoring Page**—You can configure UAC and monitor UAC authentication.
- **Allow adding a new policy and moving an existing policy to an arbitrary location**
 - **Firewall Policies Configuration Page Options**—Starting in Junos OS Release 12.1X46-D10, several new options on the Firewall Policies Configuration page are supported on high-end SRX Series devices. The Add menu includes Add before and Add after options that allow you to add a new policy before or after a selected policy. On the Move menu, there is a new Move to option that allows you to specify a target location. You can also drag and drop a policy to the target location.
 - **Checking Policies Monitoring Page**—Starting in Junos OS Release 12.1X46-D10, the Move to option on the Checking Policies Monitoring page is supported on high-end SRX Series devices.

Management Information Bases (MIBs)

- **SNMP aggregation for policy MIBs**—Starting in Junos OS Release 12.1X46-D10, this feature is supported on all SRX Series devices.

A set of systemwide policy statistics such as policy-allowed packets, bytes and rates, policy-dropped packets, bytes and rates, and policy flows allowed and rate statistics have been added in the enterprise-specific policy MIB JUNIPER-JS-POLICY-MIB. You can obtain the policy statistics by using the SNMP agent or the CLI operational mode commands. Use the following CLI commands to set, clear, and display the systemwide policy statistics:

- **set security policies policy-stats system-wide <disable | enable>**—Configures systemwide policy statistics. Disabled by default.
- **clear security policies statistics**—Clears the systemwide policy statistics.
- **show snmp mib walk jnxJsPolicySystemStats**—Displays both IPv4 and IPv6 statistics.
- **show snmp mib walk jnxJsPolicySystemStatsIPv4**—Displays only IPv4 statistics.

[See [Policy Objects MIB](#).]

Modular Interface Cards

SRX5600 and SRX5800 Services Gateway MPC Software Features—The SRX5K-MPC is a Modular Port Concentrator (MPC) that is supported on the SRX5400, SRX5600, and SRX5800.

The following features are supported on the SRX5K-MPC:

- Load balancing among SPUs using hash-based forwarding



NOTE: When the SRX5K-MPC is installed on SRX5600 and SRX5800 devices, the default session distribution mode is set to hash-based distribution mode on the devices. The hash-based distribution mode is the only mode supported on the SRX5K-MPC.

- Filtering support
- Filter-based forwarding at logical interfaces of revenue ports, firewall filter applied at loopback interface of chassis, policer applied at loopback interface of chassis
- Interface ingress policing
- Following types of threshold-based flood protection:
 - UDP-based flood protection
 - ICMP-based flood protection
 - TCP source-based SYN flood protection
 - TCP destination-based SYN flood protection

Screen

- Screen features—Starting in Junos OS Release 12.1X46-D10, the SRX5K-MPC supports screen features on the SRX5400, SRX5600, and SRX5800.

The following screen options are supported:

- Statistics-based screens
- Signature-based screens

[See [Understanding Screen Options on the SRX5000 Module Port Concentrator.](#)]

[See [Denial-of-Service Attacks Feature Guide for Security Devices.](#)]

Stream Control Transmission Protocol (SCTP)

- **SCTP payload protocol blocking**—Starting in Junos OS Release 12.1X46-D10, the **permit** traffic configuration is added to allow all types of payload protocol traffic. This feature is supported on all high-end SRX Series devices. This feature has the following enhancements:
 - The default behavior for SCTP payload protocol traffic was permit all. Now, the default behavior is drop all. However, the behavior can be changed to permit all by configuration.
 - The payload protocol traffic can be permitted by configuring the decimal value of the SCTP protocol identifiers or the name in the permit list.
 - The payload protocol traffic can be dropped by configuring the decimal value of the SCTP protocol identifiers or the name in the drop list.

[See [Understanding Stream Control Transmission Protocol.](#)]

- **Support for SCCP v20**—This feature is supported on all SRX Series devices.

Starting in Junos OS Release 12.1X46-D10, the SCCP ALG supports SCCP versions 16, 17, and 20 and several SCCP messages have been updated with a new format. Cisco Call Manager (CM) version 7 uses SCCP version 20.

[See [SCCP ALG Feature Guide for Security Devices.](#)]

- **SCTP rate limiting**—Starting in Junos OS Release 12.1X46-D10, the rate limiting functionality is extended with a generalized SCTP payload protocol rate limiting function. This feature is supported on all high-end SRX Series devices. This feature has the following enhancements:
 - The rate limiting function supports decimal identifier values for Internet Assigned Numbers Authority (IANA) SCTP protocols and synonyms for the well-known IANA SCTP protocols.
 - Each profile can be configured with many IP addresses. Each IP address can be configured with many protocols.

[See [Understanding Stream Control Transmission Protocol.](#)]

Unified Threat Management (UTM)

- **UTM antivirus, antispam, and content filtering support**—Starting in Junos OS Release 12.1X46-D10, Sophos antivirus, antispam, and content filtering features are supported on all SRX Series devices.

The existing CLI operational commands **show security utm anti-virus status** and **show security utm anti-virus statistics** have been enhanced to display the aggregated status and statistics from all Flexible PIC Concentrators (FPCs) and PICs. You can use the following new operational commands to display the status and statistics of each FPC and PIC:

- **show security utm anti-virus status fpc**
- **show security utm anti-virus status fpc fpc-slot *fpc-slot* pic-slot *pic-slot***
- **show security utm anti-virus statistics fpc**
- **show security utm anti-virus statistics fpc fpc-slot *fpc-slot* pic-slot *pic-slot***

[See [show security utm anti-virus status](#).]

[See [show security utm anti-virus statistics](#).]

- **UTM Web filtering support**—Starting in Junos OS Release 12.1X46-D10, the enhanced Web filtering feature is supported on all SRX Series devices.

The existing CLI operational commands **show security utm web-filtering status** and **show security utm web-filtering statistics** have been enhanced to display the aggregated status and statistics from all Flexible PIC Concentrators (FPCs) and PICs. You can use the following new operational commands to display the status and statistics of each FPC and PIC:

- **show security utm web-filtering status fpc**
- **show security utm web-filtering status fpc fpc-slot *fpc-slot* pic-slot *pic-slot***
- **show security utm web-filtering statistics fpc**
- **show security utm web-filtering statistics fpc fpc-slot *fpc-slot* pic-slot *pic-slot***

[See [show security utm web-filtering status](#).]

[See [show security utm anti-virus statistics](#).]

Virtual Private Networks (VPNs)

- **Enhanced X2 interface monitoring**—This feature is supported on all SRX Series devices.

In an LTE mobile network, X2 interfaces are used to connect Evolved Node Bs (eNodeBs) for signal handover, monitoring, and radio coverage. SRX Series devices connect these eNodeBs using IPsec tunnels.

This feature enables you to monitor traffic between eNodeBs by snooping into the clear text traffic as it flows from one IPsec tunnel to another. Use the **monitor-filter** statement at the **[edit security forwarding-options]** hierarchy level to duplicate clear text packets and send them to the physical interface. You can then use Ethereal or other packet analyzers to verify or collect the X2 traffic.

[See [Understanding X2 Traffic Monitoring](#) .]

- **Dead peer detection (DPD) enhancements**—This feature is supported on all SRX Series devices.

Network devices use the DPD protocol to verify the existence and availability of other peer devices. The default DPD mode **optimized** sends probes if there is no incoming IKE or IPsec traffic from the peer within a configured interval after outgoing packets are sent to the peer. The **always-send** option sends DPD probes at configured intervals regardless of traffic activity between peers. A new configuration option **probe-idle-tunnel** at the `[edit security ike gateway dead-peer-detection]` hierarchy level sends DPD probes when there is no incoming or outgoing IKE or IPsec traffic between peers.



NOTE: We recommend that you configure **probe-idle-tunnel** instead of **always-send**.

For all DPD modes, Phase 1 and Phase 2 security associations are cleared if a specified number of probes are sent with no response from the peer.

[See [Understanding Dead Peer Detection](#).]

- **IPsec VPN performance enhancements**—Starting in Junos OS Release 12.1X46-D10, a new configuration statement, **ipsec-performance-acceleration**, has been introduced under the `[edit security flow]` hierarchy to enable IPsec VPN performance acceleration. This feature is supported on SRX3400, SRX3600, SRX5600, and SRX5800 devices.

By default, VPN performance acceleration is disabled on SRX Series devices. Enabling VPN performance acceleration can improve VPN throughput under certain conditions.

The following functions are not supported:

- VPN traffic ACL accounting on physical egress and ingress interface
 - VPN traffic physical interface filter-based policer
 - VPN traffic physical interface QoS feature (classifier, remarking, scheduling, and shaping)
- **Multiple traffic selectors on a route-based VPN**—This feature is supported on all SRX Series devices.

A traffic selector (also known as a proxy ID in IKEv1) is an agreement between IKE peers to permit traffic through a tunnel if the traffic matches a specified pair of local and remote addresses. With this feature, you can define multiple traffic selectors within a specific route-based VPN, resulting in a unique SA for each traffic selector configured. Only traffic that conforms to a traffic selector is permitted through the associated IPsec SA.

To configure a traffic selector, use the **traffic-selector** configuration statement at the `[edit security ipsec vpn vpn-name]` hierarchy level. The traffic selector pair is defined with the mandatory **local-ip ip-address** and **remote-ip ip-address** statements. The CLI operational command **show security ipsec security-association traffic-selector traffic-selector** displays SA information for the specified traffic selector.

[See [Understanding Traffic Selectors in Route-Based VPNs](#).]

- **Support for IPv6 address encapsulation in route-based one-to-one site-to-site VPN tunnels**—This feature is supported on all SRX Series devices.

In tunnel mode, IPsec encapsulates the original IP datagram—including the original IP header—within a second IP datagram. The outer IP header contains the IP address of the gateway, while the inner header contains the ultimate source and destination IP addresses. The outer and inner IP headers can have a protocol field of IPv4 or IPv6. As of Junos OS Release 12.1X46-D10, the following tunnel modes are supported on SRX Series devices:

- IPv4-in-IPv4 tunnels encapsulate IPv4 packets inside IPv4 packets.
- IPv6-in-IPv6 tunnels encapsulate IPv6 packets inside IPv6 packets.
- IPv6-in-IPv4 tunnels encapsulate IPv6 packets inside IPv4 packets.
- IPv4-in-IPv6 tunnels encapsulate IPv4 packets inside IPv6 packets.

There are no new CLI configuration statements for this feature.

IPv4 and IPv6 traffic can be routed into a single IPv4 or IPv6 tunnel; the st0 interface bound to the tunnel must be configured for both family inet and family inet6. Dual stack tunnels—parallel IPv4 and IPv6 tunnels over a single physical external interface to different VPN peers—are also supported.

[See [VPN Feature Support for IPv6 Addresses](#).]

- **IKEv2 configuration payload support with RADIUS**—This feature is supported on all SRX Series devices.

Configuration payload is an Internet Key Exchange (IKE) version 2 feature used to propagate provisioning information from an IKE responder to the IKE initiator. Starting with Junos OS Release 12.1X46-D10, IKEv2 configuration payload is supported with route-based VPNs only. The following attribute types, defined in RFC 5996, *Internet Key Exchange Protocol Version 2 (IKEv2)*, can be returned to the IKE initiator by the IKE responder:

- INTERNAL_IP4_ADDRESS
- INTERNAL_IP4_NETMASK
- INTERNAL_IP4_DNS

For the IKE responder to provide the initiator with provisioning information, it must acquire the information from a specified source such as a RADIUS server. Provisioning information can also be returned from a DHCP server through a RADIUS server. On the RADIUS server, the user information should not include an authentication password. As in previous Junos OS releases for the SRX Series, the RADIUS server profile is bound to the IKE gateway using the `xauth access-profile profile-name` configuration at the `[edit security ike gateway gateway-name]` hierarchy level.

This feature is supported only for point-to-multipoint secure tunnel (st0) interfaces. For point-to-multipoint interfaces, the interfaces must be numbered and the addresses in the configuration payload INTERNAL_IP4_ADDRESS attribute type must be within the subnet range of the associated point-to-multipoint interface.



NOTE: IKEv2 on SRX Series devices does not support policy-based VPNs or VPN monitoring.

[See [Understanding Internet Key Exchange Version 2.](#)]

- **IKEv2 with NAT-T and dynamic endpoint VPN**—This feature is supported on all SRX Series devices.

Starting with Junos OS 12.1X46-D10, both IKEv2 initiators and responders in a route-based VPN can be behind NAT devices. The IKEv2 NAT-T feature supports IPsec traffic that crosses NAT devices. Static NAT and dynamic NAT are supported. In static NAT, there is a one-to-one relationship between the private and the public addresses. In dynamic NAT, there is a many-to-one or many-to-many relationship between the private and public addresses.

Dynamic endpoint (DEP) VPN is a Junos OS feature that covers IKEv2 initiator and responder perspectives. From the initiator's perspective, DEP VPN covers the situation where the IKE external interface address is not fixed and is therefore not known by the responder. This situation can occur when the peer's address is dynamically assigned by an ISP or when the peer's connection crosses a NAT device that allocates addresses from a dynamic address pool. From the responder's perspective, DEP VPN describes either a finite number of VPNs that are created for a number of VPN peers in a many-to-many scenario or a shared VPN in a many-to-one scenario.

Starting with Junos OS 12.1X46-D10, the default value for the **nat-keepalive** option configured at the `[edit security ike gateway gateway-name]` hierarchy level has been changed from 5 seconds to 20 seconds.

[See [Understanding NAT-T.](#)]

Web Authentication

- **Web-redirect firewall authentication**—Starting in Junos OS Release 12.1X46-D10, Web authentication redirect enhancement is provided on all SRX Series devices.

With this feature, when you attempt to initiate a connection across the firewall, after successful authentication the browser launches your original destination URL without your needing to retype the URL.

The following message is displayed:

Redirecting to the original url, please wait

[See [Firewall User Authentication Overview](#)]

Hardware Features


SRX5400 Services Gateway

- The SRX5400 Services Gateway (see [Figure 2 on page 111](#)) expands the SRX Series family of next-generation security platforms, delivering a high-performance, highly scalable, carrier-class security device with multiprocessor architecture. The SRX5400 Services Gateway is 5 rack units (U) tall. You can stack eight services gateways in a rack that is at least 48 U (89.3 in. or 2.24 m) in height if it has a 1 in. cap between for increased port density per unit of floor space. The services gateway provides four slots that you can populate with one Switch Control Board (SCB) and up to three additional cards comprised of an SPC and MPCs.



NOTE: The SRX5400 Services Gateway supports only the SPC II (SRX5K-SPC-4-15-320) and does not support the SRX5K-SPC-2-10-40 SPC.

Figure 2: SRX5400 Services Gateway Front Panel



NOTE: The SRX5400 Services Gateway only supports the SRX5K-MPC, and does not support older SRX5000 Series I/O cards (IOCs) or Flex IOCs cards such as:

- SRX5K-40GE-SFP
- SRX5K-4XGE-XFP
- SRX5K-FPC-IOC



NOTE: The SRX5400 Services Gateway supports Junos OS 12.1x46-D10 and later versions. It does not support previous Junos OS versions.

[See [Firewall SRX5400 Services Gateway Hardware Guide](#).]


SRX5K-MPC IOC for the SRX5000 Line of Services Gateways

The SRX5K-MPC (see [Figure 3 on page 112](#)) is an interface card with two slots that accept MICs which add Ethernet ports to your services gateway. An MPC with MICs installed functions in the same way as a regular IOC but allows you to add different types of Ethernet ports to your device. You can add just one MIC; or you can add two MICs of the same or different types.



NOTE: The SRX5K-MPC card is supported on the SRX5400, SRX5600, and SRX5800 Services Gateways. The SRX5400 Services Gateway supports only MPCs. It does not support legacy cards such as IOCs or Flex IOCs.

Figure 3: SRX5K-MPC



[See [Modular Port Concentrator SRX5K-MPC](#).]


MICs for the SRX5000 Line of Services Gateways

You use MICs and MPCs to add different combinations of Ethernet interfaces to your services gateway to suit the specific needs of your network. The following three new MICs are supported on the SRX5000 line of services gateways:

SRX-MIC-1X100G-CFP

The SRX-MIC-1X100G-CFP (see [Figure 4 on page 113](#)) can be installed in an MPC to add one 100-Gigabit Ethernet CFP port.

Figure 4: SRX-MIC-1X100G-CFP




[See [SRX-MIC-1X100G-CFP](#).]

SRX-MIC-2X40G-QSFP

The SRX-MIC-2X40G-QSFP (see [Figure 5 on page 113](#)) can be installed in an MPC to add two 40-Gigabit quad small form-factor pluggable (QSFP) Ethernet ports.

Figure 5: SRX-MIC-2X40G QSFP




[See [SRX-MIC-2X40G-QSFP](#).]

SRX-MIC-10XG-SFPP

The SRX-MIC-10XG-SFPP (see [Figure 6 on page 113](#)) can be installed in an MPC to add ten 10-Gigabit Ethernet SFP+ ports.

Figure 6: SRX-MIC-10XG SFPP



[See [SRX-MIC-10XG-SFPP](#).]

Related Documentation

- [Known Issues on page 155](#)
- [Resolved Issues on page 156](#)
- [Documentation Updates on page 183](#)
- [Changes in Behavior and Syntax on page 113](#)
- [Known Behavior on page 135](#)

Changes in Behavior and Syntax

The following current system behavior, configuration statement usage, and operational mode command usage might not yet be documented in the Junos OS documentation:

Application Firewall

- Prior to Junos OS Release 12.1X46-D10, when a rule specifies **dynamic-application junos:HTTP** without specifying any other nested application, the rule matches all HTTP traffic whether the traffic contains a nested application or not.

In Junos OS Release 12.1X46-D15 and later, that functionality has changed. When a rule specifies **dynamic-application junos:HTTP**, only HTTP traffic with no nested members is matched.

Consider the following application firewall ruleset:

```
rule-sets http-ruleset {
  rule rule1 {
    match {
      dynamic-application [junos:HTTP];
    }
    then {
      deny;
    }
  }
  default-rule {
    permit;
  }
}
```

Prior to Junos OS Release 11.4R6, the sample rules would be applied to traffic as shown in the following list:

- HTTP traffic with or without nested applications would be denied by rule1.
HTTP traffic with a nested application, such as junos:FACEBOOK or junos:TWITTER, would be denied by rule1.
- All other traffic would be permitted by the default rule.

In Junos OS Release 11.4R6 and later, the dynamic application junos:HTTP matches only the HTTP traffic that contains no recognizable nested application. The sample rules would now be applied differently:

- Only the HTTP traffic with no nested application would be denied by rule1.
HTTP traffic with a nested application, such as junos:FACEBOOK or junos:TWITTER, would no longer match rule1.
- All other traffic would be permitted by the default rule.
HTTP traffic with a nested application, such as junos:FACEBOOK or junos:TWITTER, would be permitted by the default rule.
- In Junos OS Release 12.1X46-D10 and earlier, if a nested application is not configured in any rule, then the nested application would match the default rule and take action specified in the default rule.

Starting in Junos OS Release 12.1X46-D10, the functionality has changed. If a nested application matches the default rule, then the application firewall uses the application type to match the rule and takes action specified in the rule. Use the **set security**

application-firewall nested-application dynamic-lookup enable command to control the behavior of the nested application, so that both the application and the nested application are consistent.

The default behavior of nested application before Junos OS Release 12.1X46-D10:

- Application firewall matches with the specific rule, if the nested application is configured explicitly in a rule.
- Application firewall matches with the default rule, if the nested application is not configured explicitly in a rule.
- Records the statistics of the application firewall in the matched rule.

The new behavior of nested application in Junos OS Release 12.1X46-D10:

- Application firewall matches with an application rule during application firewall policy lookup, if there is no explicit rule for the nested application.
- Application firewall matches with a specific rule, if the nested application is configured explicitly in a rule.
- Records the statistics of the application firewall in the matched rule.

Application-Level Distributed Denial of Service

- Application-level distributed denial of service, which is used to identify malicious bot clients and to drop or deny traffic if requests exceed configured thresholds, will be deprecated in future releases. As a replacement product for this feature, we recommend that you migrate to the Juniper DDoS Secure product line. For more details, contact your sales engineer.

Chassis Cluster

- In Junos OS Release 12.1X46-D10 and earlier, in a chassis cluster mode, when a secondary node failed, no notification was sent to report the secondary node failure.

Starting in Junos OS Release 12.1X46-D15, in a chassis cluster mode, when a secondary node fails, the primary node sends the SNMP trap information to report secondary node failures. New SNMP traps are added to report failures on the secondary node.

Sample SNMP trap sent when the monitored interface failed on the secondary node:

```
2014-02-18 17:36:56 10.157.83.10(via 10.157.84.10 [10.157.84.10]) TRAP, SNMP
v1, community ntrap .iso.3.6.1.4.1.2636.3.39.1.14.1 Enterprise Specific
Trap (1) Uptime: 1:29:31.53 .iso.3.6.1.4.1.2636.3.39.1.14.1.1.1.0 =
"1" .iso.3.6.1.4.1.2636.3.39.1.14.1.1.2.0 = "7"
.iso.3.6.1.4.1.2636.3.39.1.14.1.1.3.0 = "1"
.iso.3.6.1.4.1.2636.3.39.1.14.1.1.4.0 = "100"
.iso.3.6.1.4.1.2636.3.39.1.14.1.1.5.0 = "0"
.iso.3.6.1.4.1.2636.3.39.1.14.1.1.6.0 = "Priority is set to 0, Monitoring
objects are down"

2014-02-18 17:36:56 10.157.84.10 [10.157.84.10]: .iso.3.6.1.2.1.1.3.0
= Timeticks: (537153) 1:29:31.53 .iso.3.6.1.6.3.1.1.4.1.0 = OID:
.iso.3.6.1.4.1.2636.3.39.1.14.1.0.1 .iso.3.6.1.4.1.2636.3.39.1.14.1.1.1.0
= "1" .iso.3.6.1.4.1.2636.3.39.1.14.1.1.2.0 = "7"
.iso.3.6.1.4.1.2636.3.39.1.14.1.1.3.0 = "1"
```

```
.iso.3.6.1.4.1.2636.3.39.1.14.1.1.4.0 = "100"  
.iso.3.6.1.4.1.2636.3.39.1.14.1.1.5.0 = "0"  
.iso.3.6.1.4.1.2636.3.39.1.14.1.1.6.0 = "Priority is set to 0, Monitoring  
objects are down" .iso.3.6.1.6.3.1.1.4.3.0 = OID:  
.iso.3.6.1.4.1.2636.1.1.1.2.28
```

Sample SNMP trap sent when the failed interface is restored on the secondary node:

```
2014-02-18 17:38:46 10.157.83.10(via 10.157.84.10 [10.157.84.10]) TRAP, SNMP  
v1, community ntrap .iso.3.6.1.4.1.2636.3.39.1.14.1 Enterprise Specific  
Trap (1) Uptime: 1:31:20.64 .iso.3.6.1.4.1.2636.3.39.1.14.1.1.1.0 =  
"1" .iso.3.6.1.4.1.2636.3.39.1.14.1.1.2.0 = "7"  
.iso.3.6.1.4.1.2636.3.39.1.14.1.1.3.0 = "1"  
.iso.3.6.1.4.1.2636.3.39.1.14.1.1.4.0 = "0"  
.iso.3.6.1.4.1.2636.3.39.1.14.1.1.5.0 = "100"  
.iso.3.6.1.4.1.2636.3.39.1.14.1.1.6.0 = "Priority restored, Monitoring object  
failures are cleared"  
  
2014-02-18 17:38:46 10.157.84.10 [10.157.84.10]: .iso.3.6.1.2.1.1.3.0  
= Timeticks: (548064) 1:31:20.64 .iso.3.6.1.6.3.1.1.4.1.0 = OID:  
.iso.3.6.1.4.1.2636.3.39.1.14.1.0.1 .iso.3.6.1.4.1.2636.3.39.1.14.1.1.1.0  
= "1" .iso.3.6.1.4.1.2636.3.39.1.14.1.1.2.0 = "7"  
.iso.3.6.1.4.1.2636.3.39.1.14.1.1.3.0 = "1"  
.iso.3.6.1.4.1.2636.3.39.1.14.1.1.4.0 = "0"  
.iso.3.6.1.4.1.2636.3.39.1.14.1.1.5.0 = "100"  
.iso.3.6.1.4.1.2636.3.39.1.14.1.1.6.0 = "Priority restored, Monitoring object  
failures are cleared" .iso.3.6.1.6.3.1.1.4.3.0 = OID:  
.iso.3.6.1.4.1.2636.1.1.1.2.28
```

Command-Line Interface (CLI)

New or Changed CLI

- In Junos OS releases earlier than Junos OS Release 12.1X46-D25, TACACS+ options for authentication and accounting did not include an option for configuring a timestamp and time zone.

In Junos OS Release 12.1X46-D25 and later releases, you can use the **timestamp-and-timezone** option at the **[edit system tacplus-options]** hierarchy to include start time, stop time, and time zone attributes in start/stop accounting records. [See *tacplus-options*.]

- On all high-end SRX Series devices, on SPC and next-generation SPCs, IDP dedicated modes are supported only with the **inline-tap** option. In the inline-tap mode option, the **weight equal** option is not supported.

Other IDP dedicated mode configurations such as dedicated weight IDP, dedicated firewall, and dedicated equal are not supported.

The following IDP dedicated mode configuration statements are not supported:

- **set security forwarding-process application-services maximize-idp-sessions weight firewall**
- **set security forwarding-process application-services maximize-idp-sessions weight idp**

- **set security forwarding-process application-services maximize-idp-sessions weight equal**
- **set security forwarding-process application-services maximize-idp-sessions inline-tap weight equal**
- The following configuration statements are supported:
 - **set security forwarding-process application-services maximize-idp-sessions inline-tap weight firewall**
 - **set security forwarding-process application-services maximize-idp-sessions inline-tap weight idp**
- Starting in Junos OS Release 12.1X46-D10, on SRX3400 and SRX3600 devices, the value for licenses used in the output of the **show system license** command correctly displays a 1 in the full-cp-key field. Prior to this release, the output displayed a 0.
- Prior to Junos OS Release 12.1X46-D10, when you configured the DNS proxy server using the **set system services dns dns-proxy view view-name domain domain-name forwarder** CLI statement, if the IP address specified in the forwarder option was not available, the DNS query was forwarded to the default DNS servers (DNS servers provided by the ISP). The device acquired the public IP addresses from the default DNS servers.

Starting in Junos OS Release 12.1X46-D10, the **forward-only** option is added to the **set system services dns dns-proxy view view-name domain domain-name forward-only** CLI statement.

You can use the **forward-only** option to prevent the device from acquiring the public IP addresses from the DNS servers (by terminating the DNS query) in cases when the specified IP address is unreachable.

Deprecated Items for High-End SRX Series Services Gateways

[Table 12 on page 118](#) lists deprecated items (such as CLI statements, commands, options, and interfaces).

CLI statements and commands are deprecated—rather than immediately removed—to provide backward compatibility and a chance to bring your configuration into compliance with the new configuration. We strongly recommend that you phase out deprecated items and replace them with supported alternatives.

Table 12: Items Deprecated in Release 12.1

Deprecated Item	Replacement	Hierarchy Level or Command Syntax	Additional Information
download-timeout	-	download-timeout timeout	On all high-end SRX Series devices, the download-timeout command is deprecated. If the configuration is present, then the configuration is ignored. The IDP process internally triggers the security package to install when an automatic download is completed. There is no need to configure any download timeout.
node	-	request security idp security-package download	On all high-end SRX Series devices operating in a chassis cluster, the following request security idp security-package download commands with the node option are not supported: <ul style="list-style-type: none"> request security idp security-package download node primary request security idp security-package download node local request security idp security-package download node all

Table 13: Items Deprecated in Junos OS Release 12.1X46

Deprecated Item	Replacement	Hierarchy Level or Command Syntax	Additional Information
mcc-mnc	imsi-prefix	edit security gprs gtp profile profile-name apn pattern-string	On all high-end SRX Series devices, the mcc-mnc command is not supported.

Table 14 on page 119 lists the deprecated system log messages in Junos OS Release 12.1X46.

Table 14: Deprecated System Log Messages in Junos OS Release 12.1X46

Deprecated Item	Replacement
RT_GTP_PKT_ECHO_REQUEST	RT_GTP_V0_PKT_ECHO_REQUEST
RT_GTP_PKT_ECHO_REPONSE	RT_GTP_V0_PKT_ECHO_RESPONSE
	RT_GTP_V1_PKT_ECHO_REQUEST
	RT_GTP_V1_PKT_ECHO_RESPONSE
	RT_GTP_V2_PKT_ECHO_REQUEST
	RT_GTP_V2_PKT_ECHO_RESPONSE

Compatibility

- **Version compatibility for Junos SDK**—Beginning with Junos OS Release 12.1X44-D10, Junos OS applications will install on the Junos OS only if the application is built with the same release as the Junos OS release on which the application is being installed.

For example, an application built with Junos OS Release 12.1R2 will only install on Junos OS Release 12.1R2 and will not install on Junos OS Release 12.1R1 or Junos OS Release 12.1R3.

Flow and Processing

SPU software changes for the SPC—The following changes apply to all high-end SRX Series devices:

- Each SPU runs a 64-bit FreeBSD kernel instead of the 32-bit FreeBSD kernel.
- Each SPU runs a 64-bit flowd instead of the 32-bit version for increased scalability.
- With the 64-bit OS, ksynd and ifstates on the SPU run in 64-bit mode.
- **TCP initial timeout enhancement**—The minimum value you can configure for TCP session initialization is 4 seconds. The default value is 20 seconds; if required you can set the TCP session initialization value to less than 20 seconds.
- Starting with Junos OS Release 12.1X46-D10, you can configure the timeout value for a multicast flow session. In Junos OS Release 12.1X45-D10 and earlier, the timeout value for a multicast flow session was based on the packet IP protocol, which was not configurable.

Multicast flow sessions have one template flow session and one or more leaf sessions. Because these sessions are linked together, they can have only one timeout value. The earlier implementation ignored the configurable timeout values of individual policies of each leaf session, and considered only the packet IP protocol timeout, which was not configurable. For example, for UDP this timeout value was always 60 seconds. As a result, multicast streams with a packet interval of more than 60 seconds experienced premature aging-out of flow sessions and packet drops.

In the new implementation, multicast flow sessions consider the timeout values configured in leaf session policies along with the IP protocol timeout values. The highest

of these timeout values is selected as the template session timeout. You can configure the timeout value for the leaf session policy using custom applications.

[See *Configuring the Timeout Value for Multicast Flow Sessions*.]

- On all high end SRX Series devices, you can configure the TCP session timeout in a half-closed state by using the **apply-to-half-close-state** statement at the **[edit security flow tcp-session time-wait-state]** hierarchy level. This enables the system to apply the configured session timeout on receiving only one FIN packet (either client-to-server or server-to-client). When this statement is not configured, the default behavior takes effect, which is to apply the configured session timeout on receiving both the FIN packets. The default TCP session timeout remains 150 seconds. [See *apply-to-half-close-state*.]

Intrusion Detection and Prevention (IDP)

- A system log message is generated when an IDP signature database update or policy compilation fails with an empty dynamic group. The system-generated log message is **Dynamic Attack group [dyn_group_1] has no matching members found. Group is empty**.
- A new attribute, **max-synacks-queued**, is added to IDP sensor configuration TCP reassembler. This attribute defines the maximum syn/ack queued with different SEQ numbers and takes the values 0 through 5. Also, a new counter, **Duplicate Syn/Ack with different SEQ**, is added to the IDP TCP reassembler. This counter displays the number of syn/ack packets with different SEQ numbers.
- New sensor configuration options have been added to log run conditions as IDP session capacity and memory limits are approached, and to analyze traffic dropped by IDP and application identification due to exceeding these limitations.
 - At start up, traffic is ignored by IDP by default if the IDP policy is not yet loaded. The **drop-if-no-policy-loaded** option changes this behavior so that all sessions are dropped before the IDP policy is loaded.

Use the following configuration command to drop traffic before the IDP policy is loaded:

```
set security idp sensor-configuration flow drop-if-no-policy-loaded
```

The following new counters have been added to the **show security idp counters flow** command output to analyze dropped traffic due to the **drop-if-no-policy-loaded** option:

```
Sessions dropped due to no policy                                0
```

- By default, IDP ignores failover sessions in an SRX chassis cluster deployment. The **drop-on-failover** option changes this behavior and automatically drops sessions that are in the process of being inspected on the primary node when a failover to the secondary node occurs.

Use the following configuration command to drop failover sessions:

```
set security idp sensor-configuration flow drop-on-failover
```


The following new counter has been added to the **show security idp counters flow** command output to analyze dropped failover traffic due to the **drop-on-failover** option:

```
Fail-over sessions dropped                                0
```

- By default, sessions are not dropped if the IDP session limit or resource limits are exceeded. In this case, IDP and other sessions are dropped only when the device's session capacity or resources are depleted. The **drop-on-limit** option changes this behavior and drops sessions when resource limits are exceeded.

Use the following configuration commands to set or remove the **drop-on-limit** option:

```
set security idp sensor-configuration flow drop-on-limit
delete security idp sensor-configuration flow drop-on-limit
```

The following new counters have been added to the **show security idp counters flow** command output to analyze dropped IDP traffic due to the **drop-on-limit** option:

```
SM Sessions encountered memory failures                0
SM Packets on sessions with memory failures            0
SM Sessions dropped                                    0
Both directions flows ignored                          0
IDP Stream Sessions dropped due to memory failure      0
IDP Stream Sessions ignored due to memory failure      0
IDP Stream Sessions closed due to memory failure      0
Number of times Sessions exceed high mark             0
Number of times Sessions drop below low mark          0
Memory of Sessions exceeds high mark                  0
Memory of Sessions drops below low mark                0
```

The following counters have also been added to the **show security idp counters application-identification** command output to analyze dropped application identification traffic due to the **drop-on-limit** option:

```
AI-session dropped due to malloc failure before session create      0
AI-Sessions dropped due to malloc failure after create              0
AI-Packets received on sessions marked for drop due to malloc failure 0
```

The following options have been added to trigger informative log messages about current run conditions. When set, the log messages are triggered whether the **drop-on-limit** option is set or not.

- The **max-sessions-offset** option sets an offset for the maximum IDP session limit. When the number of IDP sessions exceeds the maximum session limit, a warning is logged that conditions exist where IDP sessions could be dropped. When the number of IDP sessions drops below the maximum IDP session limit minus the offset value, a message is logged that conditions have returned to normal.

```
Jul 19 04:38:13 4.0.0.254 RT_IDP: IDP_SESSION_LOG_EVENT: IDP: at 1374233893,
FPC 4 PIC 1 IDP total sessions pass through high mark 100000. IDP may drop
new sessions. Total sessions dropped 0.
```

Jul 19 04:38:21 4.0.0.254 RT_IDP: IDP_SESSION_LOG_EVENT: IDP: at 1374233901, FPC 4 PIC 1 IDP total sessions drop below low mark 99000. IDP working in normal mode. Total sessions dropped 24373.

Use the following configuration command to set the **max-sessions-offset** option:

set security idp sensor-configuration flow max-sessions-offset *offset-value*

- The **min-objcache-limit-lt** option sets a lower threshold for available cache memory. The threshold value is expressed as a percentage of available IDP cache memory. If the available cache memory drops below the lower threshold level, a message is logged stating that conditions exist where IDP sessions could be dropped because of memory allocation failures. For example, the following message shows that the IDP cache memory has dropped below the lower threshold and that a number of sessions have been dropped:

Jul 19 04:07:33 4.0.0.254 RT_IDP: IDP_SESSION_LOG_EVENT: IDP: at 1374232053, FPC 4 PIC 1 IDP total available objcache(used 4253368304, limit 7247757312) drops below low mark 3986266515. IDP may drop new sessions. Total sessions dropped 1002593.

Use the following configuration command to set the **min-objcache-limit-lt** option:

**set security idp sensor-configuration flow min-objcache-limit-lt
*lower-threshold-value***

- The **min-objcache-limit-ut** option sets an upper threshold for available cache memory. The threshold value is expressed as a percentage of available IDP cache memory. If available IDP cache memory returns to the upper threshold level, a message is logged stating that available cache memory has returned to normal. For example, the following message shows that the available IDP cache memory has increased above the upper threshold and that it is now performing normally:

Jul 19 04:13:47 4.0.0.254 RT_IDP: IDP_SESSION_LOG_EVENT: IDP: at 1374232428, FPC 4 PIC 1 IDP total available objcache(used 2782950560, limit 7247757312) increases above high mark 4348654380. IDP working in normal mode. Total sessions dropped 13424632.



NOTE: This message is triggered only if the lower threshold has been reached and the available memory has returned above the upper threshold. Fluctuations in available memory that dropped below the upper threshold but did not fall below the lower threshold would not trigger the message.

Use the following configuration commands to set the **min-objcache-limit-ut** option:

**set security idp sensor-configuration flow min-objcache-limit-ut
*upper-threshold-value***

- On all high-end SRX Series devices with a single session, when IDP is activated, the upload and download speeds are slow when compared to the firewall performance numbers.

To overcome this issue, a new CLI command, **set security idp sensor-configuration ips session-pkt-depth**, is introduced, for which the **session-pkt-depth sensor-configuration** value is global for any session.

The **session-pkt-depth sensor-configuration** value specifies the number of packets per session that are inspected by IDP. Any packets beyond the specified value are not inspected. For example, when **session-pkt-depth sensor-configuration** is configured as “n”, the IDP inspection happens only for first (n-1) packets in that session. Packets from the nth packet onwards are ignored by IDP.

The default value of **session-pkt-depth sensor-configuration** is zero. When the default value of zero is used, the session-pkt-depth value is not addressed, and IDP performs a full inspection of the session.

- Starting in Junos OS Release 12.1X46-D25, the **show security idp counters flow** command output is changed to include new fields.

[Table 3 on page 23](#) lists the output fields for the **show security idp counters flow** command. Output fields are listed in the approximate order in which they appear.

Table 15: show security idp counters flow Output Fields

Field Name	Description
Fast-path packets	Number of packets that are set through fast path after completing IDP policy lookup.
Slow-path packets	Number of packets that are sent through slow path during IDP policy lookup.
Session construction failed (Unsupported)	Number of times the packet failed to establish the session.
Session limit reached	Number of sessions that reached IDP sessions limit.
Session inspection depth reached	Number of sessions that reached inspection depth.
Memory limit reached	Number of sessions that reached memory limit.
Not a new session (Unsupported)	Number of sessions that extended beyond time limit.
Invalid index at age-out (Unsupported)	Invalid session index in session age-out message.
Packet logging	Number of packets saved for packet logging.
Policy cache hits	Number of sessions that matched policy cache.
Policy cache misses	Number of sessions that did not match policy cache.
Policy cache entries	Number of policy cache entries.

Table 15: show security idp counters flow Output Fields (*continued*)

Field Name	Description
Maximum flow hash collisions	Maximum number of packets, of one flow, that share the same hash value.
Flow hash collisions	Number of packets that share the same hash value.
Gates added	Number of gate entries added for dynamic port identification.
Gate matches (Unsupported)	Number of times a gate is matched.
Sessions deleted	Number of sessions deleted.
Sessions aged-out (Unsupported)	Number of sessions that are aged out if no traffic is received within session timeout value.
Sessions in-use while aged-out (Unsupported)	Number of sessions in use during session age-out.
TCP flows marked dead on RST/FIN	Number of sessions marked dead on TCP RST/FIN.
policy init failed	Policy initiation failed.
Number of sessions exceeds high mark	Number of sessions that exceed high mark.
Number of sessions drops below low mark	Number of sessions that fall below low mark.
Memory of sessions exceeds high mark	Session memory exceeds high mark.
Memory of sessions drops below low mark	Session memory drops below low mark.
Sessions constructed	Number of sessions established.
SM Sessions encountered memory failures	Number of SM sessions encountered memory failure.
SM Packets on sessions with memory failures	Number of SM packets on SM sessions with memory failure.
SM Sessions dropped	Number of SM sessions dropped.

Table 15: show security idp counters flow Output Fields (*continued*)

Field Name	Description
SM sessions ignored	Number of sessions ignored in Security Module (SM).
SM sessions interested	Number of SM sessions interested.
SM sessions not interested	Number of SM sessions not interested.
SM sessions interest error	Number of errors created for SM sessions interested.
Sessions destructed	Number of sessions destructed.
SM Session Create	Number of SM sessions created.
SM Packet Process	Number of packets processed from SM.
SM FTP data session ignored by IDP	Number of SM FTP data sessions that are ignored by IDP.
SM Session close	Number of SM sessions closed.
SM client-to-server packets	Number of SM client-to-server packets.
SM server-to-client packets	Number of SM server-to-client packets.
SM client-to-server L7 bytes	Number of SM client-to-server Layer 7 bytes.
SM server-to-client L7 bytes	Number of SM server-to-client Layer 7 bytes.
Client-to-server flows ignored	Number of client-to-server flow sessions that are ignored.
Server-to-client flows ignored	Number of server-to-client flow sessions that are ignored.
Both directions flows ignored	Number of server-to-client and client-to-server flow sessions that are ignored.
Fail-over sessions dropped	Number of fail-over sessions dropped.
Sessions dropped due to no policy	Number of sessions dropped because there was no active IDP policy.
IDP Stream Sessions dropped due to memory failure	Number of IDP stream sessions that are dropped because of memory failure.
IDP Stream Sessions ignored due to memory failure	Number of IDP stream sessions that are ignored because of memory failure.

Table 15: show security idp counters flow Output Fields (*continued*)

Field Name	Description
IDP Stream Sessions closed due to memory failure	Number of IDP stream sessions that are closed because of memory failure.
IDP Stream Sessions accepted	Number of IDP stream sessions that are accepted.
IDP Stream Sessions constructed	Number of IDP stream sessions that are constructed.
IDP Stream Sessions destructed	Number of IDP stream sessions that are destructed.
IDP Stream Move Data	Number of Stream data events handled by IDP.
IDP Stream Sessions ignored on JSF SSL Event	Number of IDP stream sessions that are ignored because of a JSF SSL proxy event.
IDP Stream Sessions not processed for no matching rules	Number of IDP stream sessions that are not processed for no matching rules.
IDP Stream stbuf dropped	Number of IDP stream plugin buffers dropped.
IDP Stream stbuf reinjected	Number of IDP stream plugin buffers injected.
Busy packets from stream plugin	Number of packets saved as one or more packets of this session from stream plugin.
Busy packets from packets plugin	Number of saved packets for IDP stream plugin sessions.
Bad kpp	Number of internal marked packets logged for IDP processing.
Lsys policy id lookup failed sessions	Number of sessions that failed logical systems policy lookup
Busy packets	Number of packets saved as one or more packets of this session are handed off for asynchronous processing.
Busy packet errors	Number of packets found with IP checksum error after asynchronous processing is completed.
Dropped queued packets (async mode)	Number of queued packets dropped based on policy action, reinjection failures, or if the session is marked to destruct.
Dropped queued packets failed (async mode)	Not used currently.

Table 15: show security idp counters flow Output Fields (*continued*)

Field Name	Description
Reinjected packets (async mode)	Number of packets reinjected into the queue.
Reinjected packets failed(async mode)	Number of failed reinjected packets.
AI saved processed packet	Number of AI packets saved for which the asynchronous processing is completed.
Busy packet count incremented	Number of times the busy packet count incremented in asynchronous processing.
busy packet count decremented	Number of times the busy packet count decremented in asynchronous processing.
session destructed in pme	Number of sessions destructed as a part of asynchronous result processing.
session destruct set in pme	Number of sessions set to be destructed as a result of asynchronous processing.
KQ op	Number of sessions with one of the following status: <ul style="list-style-type: none"> • KQ op hold—number of times packets held by IDP. • KQ op drop—number of times packets dropped by IDP. • KQ op route—number of times IDP decided to be route the packet directly. • KQ op Continue—number of times IDP decided to continue to process the packet. • KQ op error—number of times error occurred while IPD processing packet. • KQ op stop—number of times IDP decided to stop processing the packet.
PME wait not set	Number of AI saved packets given for signature matching.
PME wait set	Number of packets given for signature matching without AI save.
PME KQ run not called	Number of times signature matching results processed out of packet receiving order.

```
user@host> show security idp counters flow
```

IDP counter type	Value
Fast-path packets	0
Slow-path packets	0
Session construction failed	0
Session limit reached	0
Session inspection depth reached	0
Memory limit reached	0
Not a new session	0

Invalid index at ageout	0
Packet logging	0
Policy cache hits	0
Policy cache misses	0
Maximum flow hash collisions	0
Flow hash collisions	0
Gates added	0
Gate matches	0
Sessions deleted	0
Sessions aged-out	0
Sessions in-use while aged-out	0
TCP flows marked dead on RST/FIN	0
Policy init failed	0
Number of times Sessions exceed high mark	0
Number of times Sessions drop below low mark	0
Memory of Sessions exceeds high mark	0
Memory of Sessions drops below low mark	0
SM Sessions encountered memory failures	0
SM Packets on sessions with memory failures	0
Sessions constructed	0
SM Sessions ignored	0
SM Sessions dropped	0
SM Sessions interested	0
SM Sessions not interested	0
SM Sessions interest error	0
Sessions destructed	0
SM Session Create	0
SM Packet Process	0
SM ftp data session ignored by idp	0
SM Session close	0
SM Client-to-server packets	0
SM Server-to-client packets	0
SM Client-to-server L7 bytes	0
SM Server-to-client L7 bytes	0
Client-to-server flows ignored	0
Server-to-client flows ignored	0
Both directions flows ignored	0
Fail-over sessions dropped	0
Sessions dropped due to no policy	0
IDP Stream Sessions dropped due to memory failure	0
IDP Stream Sessions ignored due to memory failure	0
IDP Stream Sessions closed due to memory failure	0
IDP Stream Sessions accepted	0
IDP Stream Sessions constructed	0
IDP Stream Sessions destructed	0
IDP Stream Move Data	0
IDP Stream Sessions ignored on JSF SSL Event	0
IDP Stream Sessions not processed for no matching rules	0
IDP Stream stbuf dropped	0
IDP Stream stbuf reinjected	0
Busy pkts from stream plugin	0
Busy pkts from pkt plugin	0
bad kpp	0
Lsys policy id lookup failed sessions	0
Busy packets	0
Busy packet Errors	0
Dropped queued packets (async mode)	0
Dropped queued packets failed(async mode)	0
Reinjected packets (async mode)	0
Reinjected packets failed(async mode)	0
AI saved processed packet	0

busy packet count incremented	0
busy packet count decremented	0
session destructed in pme	0
session destruct set in pme	0
kq op hold	0
kq op drop	0
kq op route	0
kq op continue	0
kq op error	0
kq op stop	0
PME wait not set	0
PME wait set	0
PME KQ run not called	0

J-Web

- On all high-end SRX Series devices, on the Monitoring Policies page, the Deactivate and Move functions on the toolbar and the Count and Log action columns in the output table are not supported and will no longer be available.
- On all high-end SRX Series devices, on the Checking Policies page, the Delete and Deactivate buttons are not supported and will no longer be available.
- On all high-end SRX Series devices, on the Monitor > Events and Alarms > Security Events page, the *Is global policy* check box is introduced.

Logical Systems

- In Junos OS releases earlier than Junos OS Release 12.1X46-D10, when a logical tunnel interface with an IPv4 address and an Ethernet encapsulation type is configured, a configuration check is performed to ensure that the address is not identical to its peer logical tunnel interface address and that both addresses are on the same subnet. However, when a logical tunnel interface with an IPv6 address and an Ethernet encapsulation type is configured, no such configuration check is performed.

Starting in Junos OS Release 12.1X46-D10, a check is performed for IPv6 configurations. However, this change can cause existing IPv6 configurations to fail.

Management Information Bases (MIBs)

- On all high-end SRX Series devices in a chassis cluster, the calculation of the primary and secondary node sessions in the `JnxJsSPUMonitoringObjectsTable` object of the SPU monitoring MIB is incorrect. The MIB `JnxJsSPUMonitoringCurrentTotalSession` incorrectly displays total sessions.

A doubled session count is displayed because the active and backup nodes are treated as separate sessions, although these nodes are not separate sessions.

Count only the session numbers on the local node, thereby avoiding a double count, and local total sessions are displayed.

The `SPUMonitoringCurrentTotalSession` object of the MIB adds information per each SPU from the local node.

[See *SNMP MIBS and Traps Reference for SRX1400 and SRX3000 Line Services Gateways*.]

[See *SNMP MIBS and Traps Reference for SRX5000 Line Services Gateways*.]

Network Time Protocol

- When the NTP client or server is enabled in the `edit system ntp` hierarchy, the `REQ_MON_GETLIST` and `REQ_MON_GETLIST_1` control messages supported by the monlist feature within the NTP might allow remote attackers, causing a denial of service. To identify the attack, apply a firewall filter and configure the router's loopback address to allow only trusted addresses and networks.

Policy Applications

- In Junos OS releases earlier than Junos OS Release 12.1X46-D15, when you set the `count` option on a security policy using the CLI statement `security policies from-zone zone-name to-zone zone-name policy policy-name then`, the count is based on the number of packets and bytes of all network traffic that the policy allows to pass through the device.

In Junos OS Release 12.1X46-D15 and later, when you set the `count` option, the count is based on the number of packets and bytes of all network traffic the policy allows to pass through the device in both directions: the originating traffic from the client to the server (from the from-zone to the to-zone), and the return traffic from the server to the originating client.

Security Policies

- Security policies are stored in both the Routing Engine and the Packet Forwarding Engine. When you modify the policies on the Routing Engine side, the policies are synchronized to the Packet Forwarding Engine side when you commit the configuration.

The policies in the Routing Engine and Packet Forwarding Engine must always be in synchronization for the configuration to commit successfully. Under certain circumstances, policies in the Routing Engine and the Packet Forwarding Engine might be out of sync resulting in generation of system core files upon commit completion.

Starting in Junos OS Release 12.1X44-D10, the synchronization mechanism of security policies between the Routing Engine and the Packet Forwarding Engine is improved. These improvements significantly lower the probability of security policies becoming out of sync between the Routing Engine and the Packet Forwarding Engine.

However, if an out-of-sync condition occurs, the following error message will be displayed when you attempt to commit a configuration:

**Policy is out of sync between RE and PFE <SPU-name(s)>. Please resync before commit.
error: configuration check-out failed**

To re-synchronize policies between the Routing Engine and the Packet Forwarding Engine, you must:

- Reboot the device (device in standalone mode)
- Reboot both devices (devices in a chassis cluster mode)

Session Timeout for Reroute Failure

- The **route-change-timeout** configuration statement at the **[edit security flow]** hierarchy level sets the timeout when a session is rerouted but there is a reroute failure (for example, the new route uses a different egress zone from the previous route). In previous releases, the **route-change-timeout** statement was disabled by default. In Release 12.1X46-D10, the **route-change-timeout** configuration is enabled by default and the default timeout value is 6 seconds.

Simple Network Management Protocol (SNMP)

- On all high-end SRX Series devices, the screen SNMP trap **jnxJsScreenCfgChange** will not be sent during reboot.
- Prior to Junos OS Release 12.1X46-D20, the fault management system did not display the SPUs of next-generation SPCs because the XLP PICs were not defined in the MIB files. The Juniper MIBS **jnxContentsType** did not return the correct OID for next-generation SPCs.

Starting in Junos OS Release 12.1X46-D20, the **mib-jnx-chas-defines.txt** MIB file is updated with the **jnxPicType1ASPCXLP** XLP PIC. Use the **show snmp mib walk jnxContentsType** command to display the details for the XLP PIC.

Sample output displaying the incorrect OID:

```
root@host> show snmp mib walk jnxContentsType
...
```

```
jnxContentsType.8.4.1.0 = 0.0
jnxContentsType.8.4.2.0 = 0.0
jnxContentsType.8.4.3.0 = 0.0
jnxContentsType.8.4.4.0 = 0.0
...
```

For brevity, the **show** command output includes only the output that is relevant. Any other output on the system has been replaced with ellipses(...).

Sample output displaying the correct OID:

```
root@host> show snmp mib walk jnxContentsType
...
jnxContentsType.8.4.1.0 = jnxPicType1ASPCXLP
jnxContentsType.8.4.2.0 = jnxPicType2ASPCXLP
jnxContentsType.8.4.3.0 = jnxPicType2ASPCXLP
jnxContentsType.8.4.4.0 = jnxPicType2ASPCXLP
...
```

System Logs

- In Junos OS Release 12.1X46-D10 and earlier, the session-id-32 in application volume tracing (AVT) logs were not prefixed with the spu-id, whereas the flow logs were prefixed with the spu-id.

Starting in Junos OS Release 12.1X46-D10 and later, that functionality has changed. The AVT logs are now prefixed with the spu-id, so that the session-ids in AVT logs are consistent with the flow logs and unique across the system.

The following example shows session-id-32 logging before Junos OS Release 12.1X46:

```
Oct  4 09:13:14  bournville RT_FLOW: RT_FLOW_SESSION_CLOSE: session closed idle
Timeout: 4.0.0.1/9->5.0.0.1/33631 icmp 4.0.0.1/9->5.0.0.1/33631 None None 1 1
untrust trust 180000308 1(84) 0(0) 59 ICMP-ECHO UNKNOWN N/A(N/A) ge-0/0/0.0
UNKNOWN
```

```
Oct  4 09:13:14  bournville RT_FLOW: APPTRACK_SESSION_CLOSE: AppTrack session
closed idle Timeout: 4.0.0.1/9->5.0.0.1/33631 icmp ICMP-ECHO UNKNOWN
4.0.0.1/9->5.0.0.1/33631 None None 1 1 untrust trust 308 1(84) 0(0) 59 N/A N/A
No
```

The following example shows session-id-32 logging in Junos OS Release 12.1X46-D10, indicating the fix in the flow and AVT logs:

```
Oct  4 13:57:38  bournville RT_FLOW: RT_FLOW_SESSION_CREATE: session created
4.0.0.1/58565->5.0.0.1/21 junos-ftp 4.0.0.1/58565->5.0.0.1/21 None None 6 1
untrust trust 180000001 N/A(N/A) ge-0/0/0.0 UNKNOWN UNKNOWN UNKNOWN

Oct  4 13:57:38  bournville RT_FLOW: APPTRACK_SESSION_CREATE: AppTrack session
created 4.0.0.1/58565->5.0.0.1/21 junos-ftp UNKNOWN UNKNOWN
4.0.0.1/58565->5.0.0.1/21 None None 6 1 untrust trust 180000001 N/A N/A UNKNOWN
```

- On all high-end SRX Series devices, the attribute type of **packets-from-client** and **packets-from-server** options in the system logs of the following modules have been changed from unit to string:
 - AppTrack module—APPTRACK_SESSION_CLOSE, APPTRACK_SESSION_CLOSE_LS, APPTRACK_SESSION_VOL_UPDATE and APPTRACK_SESSION_VOL_UPDATE_LS
 - Session module—RT_FLOW_SESSION_CLOSE and RT_FLOW_SESSION_CLOSE_LS

On all high-end SRX Series devices, the following system log messages have been updated to include the **certificate ID**:

- PKID_PV_KEYPAIR_DEL
Existing message: **Key-Pair deletion failed**
New message: **Key-Pair deletion failed for <cert-id>**
- PKID_PV_CERT_DEL
Existing message: **Certificate deletion has occurred**
New message: **Certificate deletion has occurred for <cert-id>**
- PKID_PV_CERT_LOAD
Existing message: **Certificate has been successfully loaded**
New message: **Certificate <cert-id> has been successfully loaded**
- PKID_PV_KEYPAIR_GEN
Existing message: **Key-Pair has been generated**
New message: **Key-Pair has been generated for <cert-id>**

System Management

- On an SRX5800 device in transparent mode, if the device is not processing multicast OSPFv3 hello packets, to fix this condition you must remove the “delete security flow bridge no-packet-flooding” statement from the configuration.



NOTE: Packet flooding is enabled by default. If you have manually disabled packet flooding with the “set security flow bridge no-packet-flooding” statement, then use the configuration statement above to revert to the default behavior, which will allow the device to process multicast OSPFv3 hello packets.

- During a load override, to enhance the memory for the commit script, make sure you load the configuration by applying the following commands before commit:
set system scripts commit max-datasize 800000000
set system scripts op max-datasize 800000000

Unified Threat Management (UTM)

- Starting in Junos OS Release 12.1X46-D20, license control is supported on high-end SRX Series devices. Licensed features including anti-virus or Enhanced Web Filtering will not function until a license has been installed. The license must be installed after installing or upgrading to 12.1X46-D20. Unlicensed features such as UTM blacklists and whitelists will continue to function without a license.
- Prior to Junos OS Release 12.1X46-D20, the UTM feature profiles such as antivirus and Web filtering were provided as a default configuration regardless of the license requirement.

Starting in Junos OS Release 12.1X46-D20, the default configuration is removed. Use the **set security utm feature-profile anti-virus type <anti-virus-type>** and **set security utm feature-profile web-filtering type <web-filtering-type>** commands to configure specific antivirus and Web filter types in UTM feature profiles.

Unified In-Service Software Upgrade (ISSU)

On all high-end SRX Series devices, at the beginning of a chassis cluster unified ISSU, the system automatically fails over all RG-1+ redundancy groups that are not primary on the node from which you start the ISSU. This action ensures that the redundancy groups are all active on only the RG-0 primary node. You no longer need to fail over redundancy groups manually.

After the system fails over all RG-1+ redundancy groups, the system sets the manual failover bit and changes all RG-1+ primary node priorities to 255, regardless of whether the redundancy group failed over to the RG-0 primary node.

Virtual Private Networks (VPNs)

- For each VPN tunnel, both ESP and AH tunnel sessions are installed on SPUs and the control plane. In previous Junos OS releases, two tunnel sessions of the same protocol (ESP or AH) were installed for each VPN tunnel. For branch SRX Series devices, tunnel sessions are updated with the negotiated protocol after negotiation is completed. For high-end SRX Series devices, tunnel sessions on anchor SPUs are updated with the negotiated protocol while non-anchor SPUs retain ESP and AH tunnel sessions.

The ESP and AH tunnel sessions are displayed in the outputs for the **show security flow session** and **show security flow cp-session** operational mode commands.

- As of Junos OS Release 11.4, checks are performed to validate the IKE ID received from the VPN peer device. By default, SRX Series and J Series devices validate the IKE ID received from the peer with the IP address configured for the IKE gateway. In certain network setups, the IKE ID received from the peer (which can be an IPv4 or IPv6 address, fully qualified domain name, distinguished name, or e-mail address) does not match the IKE gateway configured on the SRX Series or J Series device. This can lead to a Phase 1 validation failure.

To modify the configuration of the SRX Series or J Series device or the peer device for the IKE ID that is used:

- Starting in Junos OS Release 12.1X46-D10, **local-address** can be configured at the **[edit security ike gateway gateway-name]** hierarchy level to specify the local gateway address when there are multiple addresses configured on an external physical interface to a VPN peer. **local-address** and the remote IKE gateway address must be in the same address family, either IPv4 or IPv6. Prior to Junos OS Release 12.1X46-D10, **local-address** was a hidden CLI configuration statement.
- On the SRX Series or J Series device, configure the **remote-identity** statement at the **[edit security ike gateway gateway-name]** hierarchy level to match the IKE ID that is received from the peer. Values can be an IPv4 or IPv6 address, fully qualified domain name, distinguished name, or e-mail address.



NOTE: If you do not configure **remote-identity**, the device uses the IPv4 or IPv6 address that corresponds to the remote peer by default.

- On the peer device, ensure that the IKE ID is the same as the **remote-identity** configured on the SRX Series or J Series device. If the peer device is an SRX Series or J Series device, configure the **local-identity** statement at the **[edit security ike gateway gateway-name]** hierarchy level. Values can be an IPv4 or IPv6 address, fully qualified domain name, distinguished name, or e-mail address.
- On all high-end SRX Series devices, the subject fields of a digital certificate can include Domain Component (DC), Common Name (CN), Organization Unit (OU), Organization (O), Location (L), State (ST), and Country (C).

In earlier releases, the **show security pki ca-certificate** and **show security pki local-certificate** CLI operational commands displayed only a single entry for each subject field, even if the certificate contained multiple entries for a field.

For example, a certificate with two OU fields such as “OU=Shipping Department, OU=Priority Mail” displayed only the first entry “OU=Shipping Department.” The **show security pki ca-certificate** and **show security pki local-certificate** CLI commands now display the entire contents of the subject field, including multiple field entries. The commands also display a new subject string output field that shows the contents of the subject field as it appears in the certificate.

- PKI objects include certificates, key pairs, and CRLs. PKI objects are read from the PKI database when the PKI Daemon starts. The PKI Daemon database loads all certificates into memory at boot time.

When an object is read into memory from the PKI database, the following new log message is created:

PKID_PV_OBJECT_READ: A PKI object was read into memory from <location>

Related Documentation

- [New and Changed Features on page 95](#)
- [Known Issues on page 155](#)
- [Resolved Issues on page 156](#)
- [Documentation Updates on page 183](#)
- [Known Behavior on page 135](#)

Known Behavior

Application Layer Gateways (ALGs)

- On all SRX Series devices, you can define the Sun RPC and MS-RPC mapping entry ageout value using the **set security alg sunrpc map-entry-timeout value** and **set security alg msrpc map-entry-timeout value** commands. The ageout value ranges from 8 hours to 72 hours, and the default value is 32 hours.

If either the Sun RPC ALG or the MS-RPC ALG service does not trigger the control negotiation even after 72 hours, the maximum RPC ALG mapping entry value times out and the new data connection to the service fails.

- The maximum size of the jbuf is 9 Kb. If the message buffer size is more than 9 Kb, the entire message cannot be transferred to the ALG packet handler. This causes subsequent packets in the session to bypass ALG handling, resulting in a transaction failure.

The limitations for SCCP ALGs are as follows:

- The SCCP is a Cisco proprietary protocol. So, any changes to the protocol by Cisco cause the SCCP ALG implementation to break. However, workarounds are provided to bypass strict decoding and allow any protocol changes to be handled gracefully.
- The SCCP ALG validates protocol data units (PDUs) with message IDs in the ranges [0x0 - 0x12], [0x20 - 0x49], and [0x81 - 0x14A]. By default, all other message IDs are treated as unknown messages and are dropped by the SCCP ALG.
- Any changes to the policies will drop the sessions and impact already established SCCP calls.
- The SCCP ALG opens pinholes that are collapsed during traffic or media inactivity. This means that during a temporary loss of connectivity, media sessions are not re-established.
- CallManager (CM) version 6.x and later does not support TCP probe packets in chassis cluster mode. As a result, the existing SCCP sessions will break when there is a failover. You can still create new SCCP sessions during failover.

The PPTP ALG with IPv6 support has the following limitation:

- Because PPP packets are compressed with Microsoft Point-to-Point Encryption (MPPE) protocol after the tunnel is set up, translation of the IP header in the PPP package cannot be handled; therefore, to make sure PPTP connection works well, the PPTP client must be able to work in dual stack mode. So that an IPv6 PPTP client can accept an IPv4 address for PPP tunnel interface, by which it can communicate with the IPv4 PPTP server without IP address translation for PPP packets.

The RTSP ALG with IPv6 support has the following limitations:

- Real-Time Streaming Protocol (RTSP) is an Application Layer protocol for controlling the delivery of data with real-time properties. The RTSP ALG supports a peer client, and the server transmits real-time media; it does not support third-party endpoints involved in the transaction.
- In case of destination NAT or NAT64 for IP address translation, if the RTSP message (including the Session Description Protocol (SDP) application content) length exceeds 2500 bytes, then the RTSP ALG processes only the first 2500 bytes of the message and ignores the rest of the message. In this scenario, the IP address in the RTSP message is not translated if the IP address does not appear in the first 2500 bytes.

The SIP ALG with IPv6 support has the following limitation:

- When NAT64 with persistent NAT is implemented, the SIP ALG adds the NAT translation to the persistent NAT binding table if NAT is configured on the Address of Record (AOR). Because persistent NAT cannot duplicate the address configured, coexistence of NAT66 and NAT64 configured on the same address is not supported.

Only one binding is created for the same source IP address.

AppSecure

- J-Web pages for AppSecure are preliminary.
- Custom application signatures and custom nested application signatures are not currently supported by J-Web.
- When ALG is enabled, application identification includes the ALG result to identify the application of the control sessions. Application firewall permits ALG data sessions whenever control sessions are permitted. If the control session is denied, there are no data sessions.

When ALG is disabled, application identification relies on its signatures to identify the application of the control and data sessions. If a signature match is not found, the application is considered unknown. Application firewall handles applications based on the application identification result.

Chassis Cluster

- If you are adding next-generation SRX5K-SPC-4-15-320 SPCs on SRX5600 and SRX5800 devices that are part of a chassis cluster, you must install the new SPCs so that a next-generation SRX5K-SPC-4-15-320 SPC is the SPC in the original lowest-numbered slot. For example, if the chassis already has two first-generation SRX5K-SPC-2-10-40 SPCs installed in slots 2 and 3, you cannot install SRX5K-SPC-4-15-320 SPCs in slot 0 or 1. You will need to make sure that an SRX5K-SPC-4-15-320 SPC is installed in the slot that provides central point functionality (in this case, slot 2). This ensures that the central point functionality is performed by an SRX5K-SPC-4-15-320 SPC.
- On all high-end SRX Series devices, IPsec VPN is not supported in active/active chassis cluster configuration (that is, when there are multiple RG1+ redundancy groups).

The following list describes the limitations for inserting an SPC on SRX1400, SRX3400, SRX3600, SRX5600, and SRX5800 devices in chassis cluster mode:

- The chassis cluster must be in active/passive mode before and during the SPC insert procedure.
- A different number of SPCs cannot be inserted in two different nodes.
- A new SPC must be inserted in a slot that is higher than the central point slot.



NOTE: The existing combo central point cannot be changed to a full central point after the new SPC is inserted.

- During an SPC insert procedure, the IKE and IPsec configurations cannot be modified.
- Users cannot specify the SPU and the IKE instance to anchor a tunnel.
- After a new SPC is inserted, the existing tunnels cannot use the processing power of the new SPC and redistribute it to the new SPC.
- Dynamic tunnels cannot load-balance across different SPCs.
- The manual VPN name and the site-to-site gateway name cannot be the same.
- In a chassis cluster scaling environment, the heartbeat-threshold must always be set to 8.
- An APN or an IMSI filter must be limited to 600 for each GTP profile. The number of filters is directly proportional to the number of IMSI prefix entries. For example, if one APN is configured with two IMSI prefix entries, then the number of filters is two.
- Eight QoS queues are supported per ae interface.
- The first recommended unified ISSU *from* release is Junos OS Release 10.4R4. If you intend to upgrade from a release earlier than Junos OS Release 10.4R4, see the release notes for the release that you are upgrading *from* for information about limitations and issues related to upgrading.
- Unified ISSU does not support UTM.
- For the latest unified ISSU support status, go to the Juniper Networks Knowledge Base (KB): <http://kb.juniper.net/> and search for KB17946.
- Unified ISSU does not support version downgrading.
- In large chassis cluster configurations on SRX1400, SRX3400 or SRX3600 devices, you need to increase the wait time before triggering failover. In a full-capacity implementation, we recommend increasing the wait to 8 seconds by modifying **heartbeat-threshold** and **heartbeat-interval** values in the **[edit chassis cluster]** hierarchy.

The product of the **heartbeat-threshold** and **heartbeat-interval** values defines the time before failover. The default values (**heartbeat-threshold** of 3 beats and **heartbeat-interval** of 1000 milliseconds) produce a wait time of 3 seconds.

To change the wait time, modify the option values so that the product equals the desired setting. For example, setting the **heartbeat-threshold** to 8 and maintaining the default value for the **heartbeat-interval** (1000 milliseconds) yields a wait time of 8 seconds. Likewise, setting the **heartbeat-threshold** to 4 and the **heartbeat-interval** to 2000 milliseconds also yields a wait time of 8 seconds.
- Packet-based forwarding for MPLS and ISO protocol families is not supported.
- On SRX5600 and SRX5800 devices, only two of the 10 ports on each PIC of 40-port 1-Gigabit Ethernet I/O cards (IOCs) can simultaneously enable IP address monitoring. Because there are four PICs per IOC, this permits a total of eight ports per IOC to be monitored. If more than two ports per PIC on 40-port 1-Gigabit Ethernet IOCs are configured for IP address monitoring, the commit will succeed but a log entry will be generated, and the accuracy and stability of IP address monitoring cannot be ensured. This limitation does not apply to any other IOCs or devices.

- IP address monitoring is not supported on reth interface link aggregation groups (LAGs) or on child interfaces of reth interface LAGs.
- Screen statistics data can be gathered on the primary device only.
- Only reth interfaces are supported for IKE external interface configuration in IPsec VPN. Other interface types can be configured, but IPsec VPN might not work.

Dynamic Host Configuration Protocol (DHCP)

- On all high-end SRX Series devices, DHCPv6 client authentication is not supported.
- On all high-end SRX Series devices, DHCP client and server functionality is not supported in a chassis cluster.
- On all high-end SRX Series devices, DHCP relay is unable to update the binding status based on DHCP_RENEW and DHCP_RELEASE messages.

Flow and Processing

- On all high-end SRX Series devices, when packet-logging functionality is configured with an improved pre-attack configuration parameter value, the resource usage increases proportionally and might affect the performance.
- On all high-end SRX Series devices, the default authentication table capacity is 45,000; the administrator can increase the capacity to a maximum of 50,000.

On SRX1400 devices, the default authentication table capacity is 10,000; the administrator can increase the capacity to a maximum of 15,000.
- On all high-end SRX Series devices, when devices are operating in flow mode, the Routing Engine side cannot detect the path MTU of an IPv6 multicast address (with a large size packet).
- On all high-end SRX Series devices, you cannot configure route policies and route patterns in the same dial plan.
- On all high-end SRX Series devices, high CPU utilization triggered for reasons such as CPU intensive commands and SNMP walks causes the BFD protocol to flap while processing large BGP updates.
- On all high-end SRX Series devices, downgrading is not supported in low-impact unified ISSU chassis cluster upgrades (LICU).
- On SRX5800 devices, network processing bundling is not supported in Layer 2 transparent mode.
- On all high-end SRX Series devices, the maximum number of concurrent sessions is 250 for SSH and Telnet, and 1024 for the Web.

General Packet Radio Service (GPRS)

The following Gateway GPRS Support Node (GGSN) and Packet Data Network Gateway (PGW) limitations are applicable for all high-end SRX Series devices.

- GGSN and PGW traffic must pass through the GPRS tunneling protocol (GTP) framework; otherwise, the tunnel status is updated incorrectly.
- The central point distributes all GTP packets to SPUs according to upstream endpoints for GGSN or PGW (one GGSN or PGW is the upstream endpoint of the GTP tunnels). Information is checked on the upstream endpoint IP and GTP packets in the GGSN pool network in the following way:
 - If the upstream endpoint source IP address in the Create-PDP-Context-Response or Create-Session-Response message is different from the IP address of the upstream endpoint, tunnels are created on one SPU. According to the IP address of the upstream endpoint for GGSN or PGW, an incoming GTP tunnel message is moved to a second SPU, and the GTP packets are dropped because no tunnel is found.



NOTE: In the GGSN pool scenario, GGSN can reply with a Create-PDP-Context-Response or Create-Session-Response message using a different source IP address than the one where the request was sent to. Therefore the request and the response can run on two different flow sessions, and these two flow sessions can be distributed to different SPUs.

The following GTP firewall limitations are applicable on all high-end SRX Series devices.

- GTP firewall does not support hot-insertable and hot-removable hardware.
- The GTP firewall needs to learn the network's GSN table and install the table for the central point and the SPU. Otherwise, some GTP traffic is blocked when the firewall is inserted in the network.
- On all high-end SRX Series devices, the GPRS tunneling protocol (GTP) module competes with other modules for memory allocation during runtime because it has dynamic memory allocation for tunnel management.
- On all high-end SRX Series devices, GTP-U inspection has the following limitations:
 - When GTP-U inspection is enabled, GTP-U throughput drops.
 - GTP-U inspection only affects the new flow sessions that are created after enabling the GTP-U inspection.



NOTE: When GTP-U inspection is disabled, the GTP module ignores the traffic on which the corresponding flow sessions were created. When GTP-U inspection is reenabled, the GTP module continues to ignore the traffic during the lifetime of the flow sessions that were created before the GTP-U inspection was reenabled.

- The ramp-up rate of GTP tunnel management messages decreases slightly (the decrease rate is less than 10 percent) when the GTP control (GTP-C) tunnel and GTP-U tunnel are created on different SPUs, whether GTP-U inspection is enabled or not.
- On all high-end SRX Series devices, NAT for GTP packets has the following limitations:
 - Only static NAT is supported; port NAT is not supported.
 - During a packet data protocol (PDP) context negotiation and update, the packet sent from the customer's GSNS must carry the public IP in the GTP payload.
 - Source IP and destination IP addresses cannot be translated simultaneously for a packet.
 - NAT for GTP only works in default logical systems.
 - IPv6 is not supported.

The following SCTP limitations are applicable on all high-end SRX Series devices:

- Dynamic policy is not supported for SCTP. You must configure all policies for needed SCTP sessions.
- SCTP modules only inspect IPv4 traffic. IPv6 traffic will be passed or dropped by flow-based or policy-based processing directly, and no SCTP module inspection will occur.
- Only the first chunk in each SCTP packet is checked.
- For static NAT to work, the interfaces packets (from one side: client or server side) coming in must belong to the same zone.
- For multihomed cases, only IPv4 Address Parameter (5) in INIT or INI-ACK is supported.
- Only static NAT is supported for SCTP.
- SCTP enable or disable is controlled by whether there is a SCTP profile configured. When you disable the SCTP feature, all associations are deleted and later SCTP packets will pass or drop according to the policy.

If you want to enable SCTP again, all the running SCTP communications will be dropped, because no associations exist. New SCTP communications can establish an association and perform the inspections.

Clear old SCTP sessions when SCTP is reenabled; doing this will avoid any impact caused by the old SCTP sessions on the new SCTP communications.

- Only established SCTP associations will be synchronized to peer node.
- A maximum of eight source IP addresses and eight destination IP addresses are allowed in an SCTP communication.
- One SPU supports a maximum of 5000 associations and a maximum of 320,000 SCTP sessions.
- The 4-way handshake process should be done in one node of a cluster. If the SCTP 4-way handshake process is handled on two nodes (for example, two sessions on two

nodes in active/active mode) or the cluster fails over before the 4-way handshake is finished, the association cannot be established successfully.

- If you configure different policies for each session belonging to one association, there will be multiple policies related to one association. The SCTP packet management (drop, rate limit, and so on) will use the profile attached to the handling SCTP session's policy.

The association's timeout will only use the profile attached to its INIT packet's policy. If the INIT packet's policy changes the attached profile, the old profile is deleted, and the association will refresh the timeout configuration. However, if the INIT packet's policy changes its attached profile without deleting the old profile, the association will not refresh the timeout configuration.

- In some cases, the associations might not be distributed to SPUs very evenly because the port's hash result on the central point is uneven. For example, this event can occur when only two peers of ports are used, and one peer has 100 associations, but another peer has only one association. In this case, the associations cannot be distributed evenly on the firewall with more than one SPU.
- SCTP sessions will not be deleted with associations, and the sessions will time out in 30 minutes, which is the default value. If you need the session to time out soon, you can preconfigure the SCTP application timeout value.
- M3UA or SCCP message parsing is checked, but the M3UA or SCCP stateful inspection is not checked.
- Only ITU-T Rec. Q.711-Q.714 (07 or 96) standard is supported. ANSI, ETSI, China, and other standards are not supported.
- Only RFC 4960 is supported.

On all high-end SRX Series devices, SCTP payload protocol blocking has the following limitations:

- The supported protocol decimal value is from 0 to 63. This value includes 48 IANA assigned protocols and 16 unassigned protocols.
- When running SCTP data traffic during a unified ISSU, the SCTP data packets are dropped at Junos OS Release 12.1X46. Only after the unified ISSU is finished, you can configure **permit** on Junos OS Release 12.1X46-D10 and pass the SCTP data traffic.
- Only the first data chunk is inspected, so protocol blocking only works for the first data chunk.

On all high-end SRX Series devices, the SCTP rate limiting function has the following limitations:

- The supported protocol decimal value is from 0 to 63. This value includes 48 IANA assigned protocols and 16 unassigned protocols.
- Only the first data chunk is inspected, so the rate limiting function only works for the first data chunk.
- A maximum of 80 addresses are rate limited in one profile.

- A maximum of 10 protocols are rate limited for one address in one profile.
- The supported rate limit value is from 1 to 12000.

Hardware

- SRX5800 devices does not support a redundant SCB card (third SCB) if an SRX5k SPC II (FRU model number: SRX5K-SPC-4-15-320) is installed on the device. If you have installed an SRX5k SPC II on an SRX5800 device with a redundant SCB card, make sure to remove the redundant SCB card.

Interfaces and Routing

This section covers filter and policing limitations.

- On SRX1400, SRX3400, and SRX3600 devices, the following feature is not supported by a simple filter:
 - Forwarding class as match condition
- On all high-end SRX Series devices, PIM does not support upstream and downstream interfaces across different virtual routers in flow mode
- On SRX1400, SRX3400 and SRX3600, devices, the following features are not supported by a policer or a three-color-policer:
 - Color-aware mode of a three-color-policer
 - Filter-specific policer
 - Forwarding class as action of a policer
 - Logical interface policer
 - Logical interface three-color policer
 - Logical interface bandwidth policer
 - Packet loss priority as action of a policer
 - Packet loss priority as action of a three-color-policer
- On all high-end SRX Series devices, the following features are not supported by a firewall filter:
 - Egress filter-based forwarding (FBF)
 - Forwarding table filter (FTF)
- SRX3400 and SRX3600 devices have the following limitations of a simple filter:
 - The forwarding class is the match condition.
 - In the packet processor on an IOC, up to 400 logical interfaces can be applied with simple filters.
 - In the packet processor on an IOC, the maximum number of terms of all simple filters is 2000.

- In the packet processor on an IOC, the maximum number of policers is 2000.
- In the packet processor on an IOC, the maximum number of three-color-policers is 2000.
- The maximum burst size of a policer or three-color-policer is 16 MB.
- On all high-end SRX Series devices, the flow monitoring version 9 has the following limitations:
 - Routing Engine based flow monitoring V5 or V8 mode is mutually exclusive with inline flow monitoring V9.
 - High-end SRX Series devices do not support multiple collectors like branch SRX Series devices. Only one V9 collector per IPv4 or IPv6 is supported.
 - Flow aggregation for V9 export is not supported.
 - Only UDP over IPv4 or IPv6 protocol can be used as the transport protocol.
 - Only the standard IPv4 or IPv6 template is supported for exporting flow monitoring records.
 - User-defined or special templates are not supported for exporting flow monitoring records.
 - Chassis cluster is supported without flow monitoring session synchronization.
- On SRX3400 and SRX3600 devices, when you enable the monitor traffic option using the **monitor traffic** command to monitor the FXP interface traffic, interface bounce occurs. You must use the **monitor traffic interface fxp0 no-promiscuous** command to avoid the issue.
- On all high-end SRX Series devices, the lo0 logical interface cannot be configured with RGO if used as an IKE gateway external interface.
- On all high-end SRX Series devices, the **set protocols bgp family inet flow** and **set routing-options flow** CLI statements are no longer available, because BGP flow spec functionality is not supported on these devices.
- On all high-end SRX Series devices, the LACP is not supported on Layer 2 interfaces.
- On all high-end SRX Series devices, BGP-based virtual private LAN service (VPLS) works on child ports and physical interfaces, but not over ae interfaces.
- When using SRX Series devices in chassis cluster mode, we recommend that you do not configure any local interfaces (or combination of local interfaces) along with redundant Ethernet interfaces.

For example:

The following configuration of chassis cluster redundant Ethernet interfaces, in which interfaces are configured as local interfaces, is not recommended:

```
ge-2/0/2 {  
  unit 0 {  
    family inet {  
      address 1.1.1/24;  
    }  
  }  
}
```



```

    }
  }

```

The following configuration of chassis cluster redundant Ethernet interfaces, in which interfaces are configured as part of redundant Ethernet interfaces, is recommended:

```

interfaces {
  ge-2/0/2 {
    gigether-options {
      redundant-parent reth2;
    }
  }
  reth2 {
    redundant-ether-options {
      redundancy-group 1;
    }
    unit 0 {
      family inet {
        address 1.1.1.1/24;
      }
    }
  }
}

```

Intrusion Detection and Prevention (IDP)

- On all high-end SRX Series devices, from Junos OS Release 11.2 and later, the IDP security package is based on the Berkeley database. Hence, when the Junos OS image is upgraded from Junos OS Release 11.1 or earlier to Junos OS Release 11.2 or later, a migration of IDP security package files needs to be performed. This is done automatically on upgrade when the IDP process comes up. Similarly, when the image is downgraded, a migration (secDb install) is automatically performed when the IDP process comes up, and previously installed database files are deleted.

However, migration is dependent on the XML files for the installed database present on the device. For first-time installation, completely updated XML files are required. If the last update on the device was an incremental update, migration might fail. In such a case, you have to manually download and install the IDP security package using the **download** or **install** CLI commands before using the IDP configuration with predefined attacks or groups.

As a workaround, use the following CLI commands to manually download the individual components of the security package from the Juniper Security Engineering portal and install the full update:

- request security idp security-package download full-update**
- request security idp security-package install**
- On all high-end SRX Series devices, the IDP policies for each user logical system are compiled together and stored on the data plane memory. To estimate adequate data plane memory for a configuration, consider these two factors:

- IDP policies applied to each user logical system are considered unique instances because the ID and zones for each user logical system are different. Estimates need to consider the combined memory requirements for all user logical systems.
- As the application database increases, compiled policies requires more memory. Memory usage should be kept below the available data plane memory to allow for database increases.
- On all high-end SRX Series devices, ingress as ge-0/0/2 and egress as ge-0/0/2.100 works with flow showing both source and destination interface as ge-0/0/2.100.
- IDP does not allow header checks for nonpacket contexts.
- On all high-end SRX Series devices, application-level distributed denial-of-service (application-level DDoS) detection does not work if two rules with different application-level DDoS applications process traffic going to a single destination application server. When setting up application-level DDoS rules, make sure that you do not configure rulebase-ddos rules that have two different application-ddos objects when the traffic destined to one application server can process more than one rule. Essentially, for each protected application server, you have to configure the application-level DDoS rules so that traffic destined for one protected server processes only one application-level DDoS rule.



NOTE: Application-level DDoS rules are terminal, which means that once traffic is processed by one rule, it will not be processed by other rules.

The following configuration options can be committed, but they will not work properly:

source-zone	destination-zone	destination-ip	service	application-ddos	Application Server
source-zone-1	dst-1	any	http	http-appddos1	1.1.1.1:80
source-zone-2	dst-1	any	http	http-appddos2	1.1.1.1:80

- On all high-end SRX Series devices, application-level DDoS rule base (rulebase-ddos) does not support port mapping. If you configure an application other than default, and if the application is from either predefined Junos OS applications or a custom application that maps an application service to a nonstandard port, application-level DDoS detection will not work.

When you configure the application setting as default, IDP uses application identification to detect applications running on standard and nonstandard ports; thus, the application-level DDoS detection would work properly.

- On all high-end SRX Series devices, all IDP policy templates are supported except All Attacks. There is a 100-MB policy size limit for integrated mode and a 150-MB policy size limit for dedicated mode. The current IDP policy templates supported are dynamic, based on the attack signatures being added. Therefore, be aware that supported templates might eventually grow past the policy size limit.

On all high-end SRX Series devices, the following IDP policies are supported:

- DMZ_Services
 - DNS_Service
 - File_Server
 - Getting_Started
 - IDP_Default
 - Recommended
 - Web_Server
- IDP deployed in both active/active and active/passive chassis clusters has the following limitations:
 - No inspection of sessions that fail over or fail back.
 - The IP action table is not synchronized across nodes.
 - The Routing Engine on the secondary node might not be able to reach networks that are reachable only through a Packet Forwarding Engine.
 - The SSL session ID cache is not synchronized across nodes. If an SSL session reuses a session ID and it happens to be processed on a node other than the one on which the session ID is cached, the SSL session cannot be decrypted and will be bypassed for IDP inspection.
 - IDP deployed in active/active chassis clusters has a limitation that for time-binding scope source traffic, if attacks from a source (with more than one destination) have active sessions distributed across nodes, then the attack might not be detected because time-binding counting has a local-node-only view. Detecting this sort of attack requires an RTO synchronization of the time-binding state that is not currently supported.

IP Monitoring

- When IP monitoring is enabled on a different subnet than the reth IP address, then you must configure the **proxy-arp unrestricted** option on the upstream router.

IPv6

- Devices with IPv6 addressing do not perform fragmentation. IPv6 hosts should either perform path MTU discovery or send packets smaller than the IPv6 minimum MTU size of 1280 bytes.
- Because IPv6 addresses are 128 bits long compared to IPv4 addresses, which are 32-bits long, IPv6 IPsec packet processing requires more resources. Therefore, a small performance degradation is observed.
- IPv6 uses more memory to set up the IPsec tunnel. Therefore, the IPsec IPv4 tunnel scalability numbers might drop.

- The addition of IPv6 capability might cause a drop in the IPsec IPv4-in-IPv4 tunnel throughput performance.
- The IPv6 IPsec VPN does not support the following functions:
 - 4in6 and 6in4 policy-based site-to-site VPN, IKE
 - 4in6 and 6in4 route-based site-to-site VPN, IKE
 - 4in6 and 6in4 policy-based site-to-site VPN, Manual Key
 - 4in6 and 6in4 route-based site-to-site VPN, Manual Key
 - 4in4, 6in6, 4in6, and 6in4 policy-based dial-up VPN, IKE
 - 4in4, 6in6, 4in6, and 6in4 policy-based dial-up VPN, Manual Key
 - Remote Access—XAuth, config mode, and shared IKE identity with mandatory XAuth
 - IKE authentication—PKI or DSA
 - IKE peer type—dynamic IP
 - Chassis cluster for basic VPN features
 - IKE authentication—PKI or RSA
 - NAT-T
 - VPN monitoring
 - Hub-and-spoke VPNs
 - NHTB
 - DPD
 - Packet reordering for IPv6 fragments over tunnels is not supported
 - Chassis cluster for advanced VPN features
 - IPv6 link-local address
- **Network and Security Manager (NSM)**—Consult the NSM release notes for version compatibility, required schema updates, platform limitations, and other specific details regarding NSM support for IPv6 addressing on all high-end SRX Series devices.
- **Security policy**—Only IDP for IPv6 sessions is supported for all high-end SRX Series devices. UTM for IPv6 sessions is not supported. If your current security policy uses rules with the IP address wildcard any, and UTM features are enabled, you will encounter configuration commit errors because UTM features do not yet support IPv6 addresses. To resolve the errors, modify the rule returning the error so that the any-ipv4 wildcard is used; and create separate rules for IPv6 traffic that do not include UTM features.

J-Web

- The following table indicates browser compatibility:

Table 16: Browser Compatibility on High-End SRX Series Devices

Device	Application	Supported Browsers	Recommended Browser
SRX1400, SRX3400, SRX3600, SRX5400, SRX5600, SRX5800	J-Web	<ul style="list-style-type: none"> • Mozilla Firefox version 3.6 or later • Microsoft Internet Explorer version 7.0 	Mozilla Firefox version 3.6 or later

- To use the Chassis View, a recent version of Adobe Flash that supports ActionScript and AJAX (Version 9) must be installed. Also note that the Chassis View is displayed by default on the Dashboard page. You can enable or disable the Chassis View using options in the dashboard Preference dialog box, but clearing cookies in Microsoft Internet Explorer also causes the Chassis View to be displayed.
- On all high-end SRX Series devices, users cannot differentiate between Active and Inactive configurations on the System Identity, Management Access, User Management, and Date & Time pages.

Logical Systems

- The master logical system must not be bound to a security profile that is configured with a 0 percent reserved CPU quota because traffic loss could occur. When upgrading all high-end SRX Series devices from Junos OS Release 11.2, make sure that the reserved CPU quota in the security profile that is bound to the master logical system is configured for 1 percent or more. After upgrading from Junos OS Release 11.2, the reserved CPU quota is added to the default security profile with a value of 1 percent.
- On all high-end SRX Series devices, quality-of-service (QoS) classification across interconnected logical systems does not work.
- On all high-end SRX Series devices, the number of logical system security profiles you can create is constrained by an internal limit on security profile IDs. The security profile ID range is from 1 through 32, with ID 0 reserved for the internally configured default security profile. When the maximum number of security profiles is reached, if you want to add a new security profile, you must first delete one or more existing security profiles, commit the configuration, and then create the new security profile and commit it. You cannot add a new security profile and remove an existing one within a single configuration commit.

If you want to add more than one new security profile, the same rule is true. You must first delete the equivalent number of existing security profiles, commit the configuration, and then create the new security profiles and commit them.

- **User and administrator configuration for logical systems**—Configuration for users for all logical systems and all user logical systems administrators must be done at the root level by the master administrator. A user logical system administrator cannot create other user logical system administrators or user accounts for their logical systems.

- **Name-space separation**—The same name cannot be used in two logical systems. For example, if logical-system1 includes the username “Bob” then other logical systems on the device cannot include the username “Bob”.
- **Commit rollback**—Commit rollback is supported at the root level only.
- **Trace and debug**—Trace and debug are supported at the root level only.
- **Class of service**—You cannot configure class of service on logical tunnel (lt-0/0/0) interfaces.
- **ALGs**—The master administrator can configure ALGs at the root level. The configuration is inherited by all user logical systems. It cannot be configured discretely for user logical systems.

Network Address Translation (NAT)

- **Single IP address in a source NAT pool without PAT**—The number of hosts that a source NAT pool without PAT can support is limited to the number of addresses in the pool. When you have a pool with a single IP address, only one host can be supported, and traffic from other hosts is blocked because there are no resources available.

If a single IP address is configured for a source NAT pool without PAT when NAT resource assignment is not in active-backup mode in a chassis cluster, traffic through node 1 will be blocked.

- For all ALG traffic, except FTP, we recommend that you not use the static NAT rule options **source-address** or **source-port**. Data session creation can fail if these options are used, because the IP address and the source port value, which is a random value, might not match the static NAT rule. For the same reason, we also recommend that you not use the source NAT rule option **source-port** for ALG traffic.

For FTP ALG traffic, the **source-address** option can be used because an IP address can be provided to match the source address of a static NAT rule.

Additionally, because static NAT rules do not support overlapping addresses and ports, they should not be used to map one external IP address to multiple internal IP addresses for ALG traffic. For example, if different sites want to access two different FTP servers, the internal FTP servers should be mapped to two different external IP addresses.

- On all high-end SRX Series devices, in case of SSL proxy, sessions are whitelisted based on the actual IP address and not on the translated IP address. Because of this, in the whitelist configuration of the SSL proxy profile, the actual IP address should be provided and not the translated IP addresses.

Example:

Consider a destination NAT rule that translates destination IP address 20.20.20.20 to 5.0.0.1 using the following commands:

- **set security nat destination pool d1 address 5.0.0.1/32**
- **set security nat destination rule-set dst-nat rule r1 match destination-address 20.20.20.20/32**
- **set security nat destination rule-set dst-nat rule r1 then destination-nat pool d1**

In the above scenario, to exempt a session from SSL proxy inspection, the following IP address should be added to the whitelist:

- **set security address-book global address ssl-proxy-exempted-addr 20.20.20.20/32**
- **set services ssl proxy profile ssl-inspect-profile whitelist ssl-proxy-exempted-addr**
- Maximum capacities for source pools and IP addresses have been extended on all high-end SRX Series devices as follows:

Pool/PAT Maximum Address Capacity	SRX1400	SRX3400 SRX3600	SRX5400 SRX5600 SRX5800
Source NAT pools	8192	8192	12,288
IP addresses supporting port translation	8192	8192	12,288
PAT port number	256M	256M	384M

Increasing the capacity of source NAT pools consumes memory needed for port allocation. When source NAT pool and IP address limits are reached, port ranges should be reassigned. That is, the number of ports for each IP address should be decreased when the number of IP addresses and source NAT pools is increased. This ensures NAT does not consume too much memory. Use the **port-range** statement in configuration mode in the CLI to assign a new port range or the **pool-default-port-range** statement to override the specified default.

Configuring port overloading should also be done carefully when source NAT pools are increased.

For source pool with PAT in range (63,488 through 65,535), two ports are allocated at one time for RTP or RTCP applications, such as SIP, H.323, and RTSP. In these scenarios, each IP address supports PAT, occupying 2048 ports (63,488 through 65,535) for ALG module use. On SRX5600 and SRX5800 devices, if all of the 12288 source pool is configured, a port allocation of 2M is reserved for twin port use.

- **NAT rule capacity change**—To support the use of large-scale NAT at the edge of the carrier network, the device wide NAT rule capacity has been changed.

The number of destination, static, and source NAT rules has been incremented as shown in [Table 17 on page 152](#). The limitation on the number of destination rule sets and static rule sets has been increased.

[Table 17 on page 152](#) provides the requirements per device to increase the configuration limitation as well as to scale the capacity for each device.

Table 17: Number of Rules on All High-End SRX Series Devices

NAT Rule Type	SRX1400	SRX3400 SRX3600	SRX5400 SRX5600 SRX5800
Source NAT rule	8192	20480	30720
Destination NAT rule	8192	20480	30720
Static NAT rule	8192	20480	30720

The restriction on the number of rules per rule set has been increased so that there is only a devicewide limitation on how many rules a device can support. This restriction is provided to help you better plan and configure the NAT rules for the device.

For memory consumption, there is no guarantee to support these numbers (maximum source rule or rule set + maximum destination rule or rule set + maximum static rule or rule-set) at the same time for SRX1400, SRX3400, SRX3600, SRX5600, and SRX5800 devices.

The suggested total number of rules and rule sets is listed in following table:

Objects	SRX3400 SRX3600	SRX5400 SRX5600 SRX5800
Total NAT rule sets per system	20,480	30,720
Total NAT rules per rule set	20,480	30,720

Security Policies

- On all high-end SRX Series devices, the current SSL proxy implementation has the following connectivity limitations:
 - The SSLv2 protocol is not supported. SSL sessions using SSLv2 are dropped.
 - SSL sessions where client certificate authentication is mandatory are dropped.
 - SSL sessions where renegotiation is requested are dropped.
- On all high-end SRX Series devices, for a particular session, the SSL proxy is only enabled if a relevant feature related to SSL traffic is also enabled. Features that are related to SSL traffic are IDP, application identification, application firewall, and application tracking. If none of the above listed features are active on a session, the SSL proxy bypasses the session and logs are not generated in this scenario.
- On all high-end SRX Series devices, you cannot configure the following IP addresses as negated addresses in a policy:
 - Wildcard addresses
 - IPv6 addresses

- Addresses such as **any**, **any-ipv4**, **any-IPv6** and **0.0.0.0**
- When a range of addresses or a single address is negated, it can be divided into multiple addresses. These negated addresses are shown as a prefix or a length that requires more memory for storage on a Packet Forwarding Engine.
- Each platform has a limited number of policies with negated addresses. A policy can contain 10 source or destination addresses. The capacity of the policy depends on the maximum number of policies that the platform supports.

Services Offloading

- Services offloading has the following limitations:
 - Transparent mode is not supported. If transparent mode is configured, a normal session is installed.
 - LAG is not supported. If a LAG is configured, a normal session is installed.
 - Only multicast sessions with one fan-out are supported. If a multicast session with more than one fan-out exists, a normal session is installed.
 - Only active/passive chassis cluster configuration is supported. Active/active chassis cluster configuration is not supported.
 - Fragmented packets are not supported. If fragmented packets exist, a normal session is installed.
 - IPv6 is not supported. If IPv6 is configured, a normal session is installed.



NOTE: A normal session forwards packets from the network processor to the SPU for fast-path processing. A services-offload session processes fast-path packets in the network processor and the packets exit out of the network processor itself.

- For Non-Services-Offload Sessions:
 - When services offloading is enabled, for normal sessions, the performance can drop by approximately 20 percent for connections per second (CPS) and 15 percent for packets per second (pps) when compared with non-services-offload mode.
 - For Services-Offload Sessions:

When services offloading is enabled, for fast-forward sessions, the performance can drop by approximately 13 percent for connections per second (CPS).

Simple Network Management Protocol (SNMP)

- On all high-end SRX Series devices, the **show snmp mib** CLI command will not display the output for security related MIBs. We recommend that you use an SNMP client and prefix **logical-system-name@** to the community name. For example, if the community is **public**, use **default@public** for default root logical system.

Unified Access Control

- During SRX device communication to the Infranet Controller (IC), the connection remains in attempt-next state preventing a successful communication. This happens when an outgoing interface used to connect the IC is a part of routing-instance.

Unified Threat Management (UTM)

- On SRX5400 devices configured with Sophos Antivirus, certain files whose sizes are larger than the max-content-size might not go into fallback unlike other AV engines and instead end up being detected as clean file for few protocols which does not pre-declare the content size.

Virtual Private Networks (VPNs)

On all high-end SRX Series devices, IKEv2 does not include support for:

- Policy-based tunnels
- Dial-up tunnels
- VPN monitoring
- NHTP for st0—Reusing the same tunnel interface for multiple tunnels
- EAP
- Multiple child SAs for the same traffic selectors for each QoS value
- Proposal enhancement features
- Reuse of Diffie-Hellman (DH) exponentials
- IP Payload Compression Protocol (IPComp)
- On all high-end SRX Series devices, configuring XAuth with AutoVPN secure tunnel (st0) interfaces in point-to-multipoint mode and dynamic IKE gateways is not supported.
- VPN monitoring and Suite B cryptographic configuration options **ecdsa-signatures-384** (for IKE authentication) and Diffie-Hellman **group20** consume considerable CPU resources. If VPN monitoring and the **ecdsa-signatures-384** and **group20** options are used on an SRX Series device with a large number of tunnels configured, the device must have the next-generation SPC installed.
- On all high-end SRX Series devices, for auto VPN, the tunnel setup rate decreases with an increase in the number of SPCs in the device.

- IPv6 policy-based VPN is not supported
- On all high-end SRX Series devices, DH-group 14 is not supported for dynamic VPN.
- On all high-end SRX Series devices, when you enable VPN, overlapping of the IP addresses across virtual routers is supported with the following limitations:
 - An IKE external interface address cannot overlap with any other virtual router.
 - An internal or trust interface address can overlap across any other virtual router.
 - An st0 interface address cannot overlap in route-based VPN in point-to-multipoint tunnels such as NHTB.
 - An st0 interface address can overlap in route-based VPN in point-to-point tunnels.
- On all high-end SRX Series devices, the DF-bit configuration for VPN only works if the original packet size is smaller than the st0 interface MTU, and larger than the **external interface-ipsec overhead**.
- RIP is not supported in point-to-multipoint (P2MP) VPN scenarios including AutoVPN deployments. We recommend OSPF or IBGP for dynamic routing when using P2MP VPN tunnels.
- On all high-end SRX Series devices, the IPsec NAT-T tunnel scaling and sustaining issues are as follows:
 - For a given private IP address, the NAT device should translate both 500 and 4500 private ports to the same public IP address.
 - The total number of tunnels from a given public translated IP cannot exceed 1000 tunnels.
- On all high-end SRX Series devices, for auto VPN, the tunnel setup rate decreases with an increase in the number of SPCs in the device.

Related Documentation

- [New and Changed Features on page 95](#)
- [Resolved Issues on page 156](#)
- [Known Issues on page 155](#)
- [Documentation Updates on page 183](#)
- [Changes in Behavior and Syntax on page 113](#)

Known Issues

The following problems currently exist in Juniper Networks SRX Series Services Gateways. The identifier following the descriptions is the tracking number in the Juniper Networks Problem Report (PR) tracking system.



NOTE: For the latest, most complete information about outstanding and resolved issues with the Junos OS software, see the Juniper Networks online software defect search application at <http://www.juniper.net/prsearch>.

Intrusion Detection and Prevention (IDP)

- On all high-end SRX Series devices, an active IDP session might not be fetched by the **show security flow session idp** and **show security flow session summary idp** commands. This issue occurs only when both IDP and AppID are enabled. [PR1045587](#)

Network Address Translation (NAT)

- On all high-end SRX Series devices, when source NAT is configured, the ports are allocated randomly by default. In rare circumstances, the global random port table of source pools or interfaces becomes damaged by certain services or traffic. This damage can result in low-range ports being assigned a higher priority in sessions. Ports might be reused quickly, causing application access failure. [PR1006649](#)

Platform and Infrastructure

- On all high-end SRX Series devices, when you connect to the device through wireless AP the secure access port incorrectly allows access to the MAC addresses that are not in the list of allowed MAC addresses. [PR587163](#)

Unified Threat Management (UTM)

- On SRX3400 and SRX3600 devices, when you enable UTM and if IPv6 is configured on reth interface running LACP, Duplicated Address Detection (DAD) might fail under certain conditions.

As a workaround, disable DAD using the **set interface reth[x] unit [y] family inet dad-disable** command. [PR1012122](#)

Virtual Private Networks (VPN)

- On all high-end SRX Series devices with policy-based IPsec VPN configured, deleting security policies that are associated with a VPN tunnel might result in a stale VPN tunnel remaining. In addition, the tunnel might be associated with the newly added security policies. [PR1034049](#)

Related Documentation

- [New and Changed Features on page 95](#)
- [Known Behavior on page 135](#)
- [Documentation Updates on page 183](#)

Resolved Issues

The following are the issues that have been resolved in Junos OS Release 12.1X46 for Juniper Networks SRX Series Services Gateways. The identifier following the description is the tracking number in the Juniper Networks Problem Report (PR) tracking system.



NOTE: For the latest, most complete information about outstanding and resolved issues with the Junos OS software, see the Juniper Networks online software defect search application at <http://www.juniper.net/prsearch>.

Resolved Issues - 12.1X46-D35

Application Identification

- On all high-end SRX Series devices running Junos OS Release 12.1X46 and earlier, if application identification (AppID) is enabled, performance degradation is seen in comparison with devices running Junos OS Release 12.1X47-D10 and later. This is because the AppID function does not ignore the related sessions when AppID has reached the terminal state, and continues with the serialization processing for those sessions. It is important to note that Junos OS Release 12.1X47 and later releases use advanced AppID. [PR1046509](#)

Application Layer Gateways (ALGs)

- On all high-end SRX Series devices in a chassis cluster with TCP-based ALG enabled and the TCP keepalive mechanism used on the TCP server and client, after a data plane redundancy group (RG1+) failover, the keepalive message causes the mbuf to be held by the ALG until the session timeout. As a result, a high mbuf usage alarm is generated. Application communication failure occurs due to lack of mbuf. [PR1031910](#)
- On all high-end SRX Series devices with the SIP ALG and NAT enabled, if you place a call on hold or off hold many times, each time with different media ports, the resource in the call is used, resulting in one-way audio. Tearing down the call clears the resource, and following calls are not affected. [PR1032528](#)
- On all high-end SRX Series devices in a chassis cluster with the SCCP ALG enabled and if the SCCP state in use flag is not configured in the process of the SCCP call in the device, the related real-time object (RTO) hot synchronization might cause the flowd process to crash. [PR1034722](#)
- On all high-end SRX Series devices with the MS-RPC ALG enabled, the flowd process might crash when the MS-RPC ALG processes the crafted **ISystemActivator RemoteCreateInstance Response** packets. [PR1036574](#)
- On all high-end SRX Series devices with the SIP ALG and NAT enabled, the SIP ALG does not execute IP translation for the retransmitted 183 session progress messages. In this scenario, the SIP call will fail when the device receives the first 183 session progress messages without SDP information, but the retransmitted 183 session progress messages contains SDP information. [PR1036650](#)
- On all high-end SRX Series devices, the DNS ALG does not terminate the session when a truncated DNS reply is received. Hence, the session remains up until high timeout (10~50) is reached. [PR1038800](#)
- On all high-end SRX Series devices, the **SIP ALG decode packet** error occurs in the system log when the unsupported blank packets are used as keepalive messages. [PR1057170](#)

- On all high-end SRX Series devices, the current SIP parser does not parse the quotation marks in the mime message boundary, and the message body of the SIP messages might be cut off. [PR1064869](#)
- On all high-end SRX Series devices with the MS-RPC ALG enabled, the flowd process might crash due to incorrect MS-RPC ALG parsing for the **ISystemActivator RemoteCreateInstance Response** packets. [PR1066697](#)

Authentication

- On all high-end SRX Series devices with firewall authentication configured, an authentication entry leak on the data plane occurs when an authenticated user tries to re-authenticate. As a result, firewall authentication will not allow anymore authentication entries to be created. [PR969085](#)
- On all high-end SRX Series devices with firewall authentication enabled, when a firewall authentication from an authenticated IP address for a new authentication fails, and then a pass-through firewall authentication tries this entry, the firewall authentication function accesses a freed memory, which results in a flowd process crash. [PR1040214](#)
- On all high-end SRX Series devices with firewall authentication enabled, in a rare timing condition, if there are many pending sessions in a firewall authentication entry with failed state, then a packet entering and matching this failed authentication entry might cause the flowd process to crash. [PR1048623](#)

Chassis Cluster

- On SRX5400, SRX5600, and SRX5800 devices with SPC II cards installed, when IP spoofing is enabled, after the device under test (DUT) is rebooted, the address books in the Packet Forwarding Engine will be removed and not pushed back into the Packet Forwarding Engine. Due to this issue, IP spoofing does not work after the reboot. [PR920216](#)
- On all high-end SRX Series devices in chassis cluster mode, during control plane RGO failover, a policy resynchronization operation compares the policy message between the Routing Engine and the Packet Forwarding Engine. However, some fields in the security policy data message are not processed. Data for unprocessed fields might be treated differently and cause the flowd process to crash. [PR1040819](#)

Dynamic Host Configuration Protocol (DHCP)

- On all high-end SRX Series devices configured as a DHCP server (using the `jdhcpd` process), when the DHCP server gets a new request from a client and applies an IP address from the authentication process (`authd`), the `jdhcpd` process communicates with `authd` process twice as expected (once for the DHCP discovery message and once for the DHCP request message). If the authentication fails in the first message, the `authd` process will indefinitely wait for the second authentication request. However, the `jdhcpd` process never sends the second request, because the process detects that the first authentication did not occur. This causes memory leak on the `authd` process, and the memory might get exhausted, generating a core file and preventing DHCP server service. High CPU usage on the Routing Engine might also be observed. [PR1042818](#)

General Packet Radio Service (GPRS)

- On all high-end SRX Series devices in a mobile packet core network, with GTPv2 enabled and the device configured as a border gateway, the GTP packets might be dropped with a missing information element drop reason message. The packets are dropped because the information element check in processing the GTPv2 modify bearer request is not accurate. The check should only exist when Tracking Area Updates (TAU), Routing Area Updates (RAU), or handover are processed with a Serving Gateway (SGW) change on the S5/8 interface. [PR1065958](#)

Flow-Based and Packet-Based Processing

- On all high-end SRX Series devices, the Network Processing Unit (NPU) TCP sequence check might cause TCP packets to be dropped if the services-offload policy permits. [PR891118](#)
- On all high-end SRX Series devices, when composite next hop is used, RSVP session flap might cause an if state mismatch between the master Routing Engine and the backup Routing Engine, leading to a kernel crash on the master Routing Engine. [PR905317](#)
- On all high-end SRX Series devices with IDP configured, in rare cases, where the device runs out of memory, the flowd process might crash if shell code detection occurs. [PR985139](#)
- On all high-end SRX Series devices, when you configure **http-get RPM probes** to measure the website response, the probes might fail because the HTTP server might incorrectly interpret the request coming from the device. [PR1001813](#)
- On all multiple thread-based high-end SRX Series devices, if IDP, AppSecure, ALG, GTP, or the SCTP feature, which is required for serialization flow processing is enabled, the device might encounter an issue where two flow threads work on the same session at the same time for the serialization flow processing. This issue might cause memory corruption, and then result in a flowd process crash. [PR1026692](#)
- On all high-end SRX Series devices, when a device forwards traffic, a flowd core file is generated. This is a generic issue and does not impact any feature. [PR1027306](#)
- On SRX3400, SRX3600, SRX5400, SRX5600, and SRX5800 devices, if session distribution mode is set to hash mode, the TCP connections that are required for session

serialization processing might not be established. This is because of incorrect processing of the 3-way handshake in the TCP proxy module. [PR1037822](#)

- On all high-end SRX Series devices in a chassis cluster Z mode, if static NAT or destination NAT is configured, and in the NAT rule, the IP address of the incoming interface is used as a matching condition of the destination address (for example, **set security nat static <rule-set-name> rule <rule-name> match destination-address <use the IP address of incoming interface>**), then the traffic matching the NAT rule is discarded. [PR1040185](#)
- On all high-end SRX Series devices with GRE tunnel configured, the carrier interface of GRE tunnel is not updated when a more accurate and new route to the tunnel destination address is added, which might cause traffic loss in some scenarios. [PR1040666](#)
- On all high end SRX Series devices, when self-generated traffic is processed by IDP, the IDP function might trigger an unmatched flow lock operation, which leads to a dead lock condition, and eventually causes the flowd process to crash. [PR1046801](#)
- On all high-end SRX Series devices in transparent mode, when the **PIM register-stop** message passes through the device, the device cannot match the PIM session that is created by the register packet. The **PIM register-stop** message tries to create a new session, and the session is dropped during the session creation process due to a session conflict. [PR1049946](#)
- On all high-end SRX Series devices running Junos OS Release 12.3X48-D10 or later, with enhanced Web filtering configured, the connection to the Websense ThreatSeeker Intelligence Cloud might time out if **strict-syn-check** is enabled under the **[security flow tcp-session]** hierarchy. [PR1061064](#)

Hardware

- On all high-end SRX Series devices, due to an I2C hardware issue, the power entry module (PEM) status register is unstable. As a result, the chassisd reports the wrong power state. [PR1047547](#)

Interfaces and Routing

- On all high-end SRX Series devices, when a router is acting as an NTP broadcast server, broadcast addresses must be in the default routing instance. NTP messages are not broadcasted when the address is configured in a VPN virtual routing and forwarding (VRF) instance. [PR887646](#)
- On all high-end SRX Series devices, LAG interface gratuitous ARP is neither generated nor sent out on the link when **gratuitous-arp-on-ifup** is configured. [PR889851](#)
- On all high-end SRX Series devices, the **clear security dns-cache** command is extended to resolve all DNS entries immediately. Similarly, the security policies containing DNS names are updated immediately to use the refreshed IP addresses after the FQDN addresses are resolved. [PR970235](#)
- On all high-end SRX Series devices, during the ISSU process, the Packet Forwarding Engine connects and sometimes disconnects the Routine Engine. Hence, the IP resolve events sent to the Packet Forwarding Engine are ignored. When you configure multiple

DNS policies after the ISSU process, some of the policies will not have IP addresses in the Packet Forwarding Engine. [PR985731](#)

- On all high-end SRX Series devices, the **commit synchronize** command fails because the kernel socket gets stuck. [PR1027898](#)
- On SRX1400, SRX3400, and SRX3600 devices, memory leak occurs on the Control Plane Processor (CPP) logical interfaces are deleted and the interprocess communication messages are received by the CPP. High memory usage on the CPP might be seen in an interface flapping situation. [PR1059127](#)

J-Web

- On all high-end SRX Series devices, if a security policy contains a **tcp-options** statement, modifying this security policy by using J-Web results in the loss of the **tcp-options** statement. This is because the **tcp-options** configuration is missing in the J-Web security policy configuration. [PR1063593](#)

MIBs

- On all high-end SRX Series devices, there are compilation issues with the `mib-jnx-license`, `mib-jnx-sp-nat`, and `mib-jnx-subscriber` MIBs. [PR794327](#)

Network Address Translation (NAT)

- On all high-end SRX Series devices with persistent NAT enabled, if an invalid flow with the protocol value 0 creates a persistent NAT entry, then this persistent NAT entry is not cleared even when the invalid session is cleared. [PR935325](#)
- On all high-end SRX Series devices configured in chassis cluster mode, when ALG traffic performs NAT translation, in a rare condition, invalid ALG binding entries might be created on the secondary node, which results in a flowd process crash on the secondary node. [PR1037617](#)

Platform and Infrastructure

- On all high-end SRX Series devices, the packets per second (pps) and bits per second (bps) counters are not reporting accurate values while checking the **monitor traffic interface interface-name** command or the **show interface interface-name extensive** command. [PR1033222](#)
- On all high-end SRX Series devices, the configurations of group `junos-defaults` are lost after a configuration rollback. As a result, the **commit** command fails. [PR1052925](#)

Security Policy

- On all high-end SRX Series devices, when two security policies are combined and the whole address space is used, then the secondary security policy might fail to evaluate traffic. [PR1052426](#)
- On all high-end SRX Series devices, changing a dynamic address of a security policy might cause its dynamic address identification to be mismatched between the Routing Engine and the Packet Forwarding Engine due to the difference between the new and the old configuration being ignored. [PR1061253](#)

- On all high-end SRX Series devices configured in a chassis cluster, the **count** option in security policy might not work after failover. This is because the Packet Forwarding Engine does not resend the message with policy states to the Routing Engine after failover. The policy lookup counter might disappear when you execute the **show security policies from-zone * to-zone * policy-name * detail | grep lookups** command. [PR1063654](#)

Unified Threat Management (UTM)

- On all high-end SRX Series devices, when UTM Sophos antivirus is enabled and a file that is not supported by Sophos antivirus is transferred through SMTP, the device might not be able to handle the last packet, and mail will be on hold. When packets are later sent on this session, the packet that was on hold will be handled by the device and the system will return to normal state. [PR1049506](#)
- On all high-end SRX Series devices, if the name server is configured and the interface pointing to the name server is down, in a rare condition, the flowd process might crash due to a UTM internal function even though UTM is not configured. [PR1066510](#)

Virtual Private Networks (VPNs)

- On all high-end SRX Series devices with IPsec VPN configured using IKE version 1, the device can hold only two pairs of IPsec security associations (SAs) per tunnel. When the third IPsec SA rekey occurs, the oldest IPsec SA is deleted. Due to this mechanism, a looping of IPsec SA rekey might occur. For example, when a VPN peer contains incorrect configuration that has more than two proxy IDs matching only one proxy ID on a device, the rekey looping issue might cause the flowd process to crash on multiple thread-based SRX Series devices. [PR996429](#)
- On all high-end SRX Series devices, in a hub-and-spoke IPsec VPN scenario, on the hub site, when committing the static NHTBs on the multipoint secure tunnel (st0) interface, the VPN routes might become active even though the VPN tunnel is down. This issue also occurs when the system reboots with static NHTBs and the related static routes configured. [PR1007235](#)
- On all high-end SRX Series devices, the block size for Advanced Encryption Standard (AES) in Galois/Counter Mode (GCM) has been reduced from 8 to 4. Block size 8 is used for connecting to other SRX Series devices, and block size 4 is interoperable with systems from Cisco, strongSwan, and other companies. When you set the correct block size 4 for AES-GCM, it causes a problem when connecting to previous releases of Junos OS for SRX Series devices. The problem affects certain packet sizes, so it might appear to work for some traffic, such as ping, but not for other traffic. In a hub-and-spoke configuration, the upgrade causes problems with tunnels to all spokes until they are upgraded. [PR1037432](#)
- On all high-end SRX Series devices, when a primary IP address of an interface changes, some IPsec tunnels terminated on that interface might go down. [PR1044620](#)
- On all high-end SRX Series devices configured with a large number of IPsec VPN tunnels, in a very rare condition, if VPN monitoring is enabled, the kmd process might crash when you delete the partial VPN tunnels. [PR1044660](#)
- On all high-end SRX series devices, in a tunnel over route-based IPsec VPN, GRE or IP-in-IP tunnel scenario, such as IPsec VPN over GRE tunnel, after the encapsulation

of the first tunnel, the next-hop in internal processing might not be set properly to point to the second tunnel, which results in packet loss. [PR1051541](#)

Resolved Issues - 12.1X46-D30

Application Layer Gateways (ALGs)

- On all high-end SRX Series devices, when the ALG processes SIP traffic, a memory corruption issue might occur, causing the flowd process to crash. [PR992478](#)
- On all high-end SRX Series devices with the MS-RPC ALG enabled, occasionally, when more than one IP and port pair exist in the MS RPC response packet, and if these IP and port pair are the same, the ALG group might leak. This issue might occur even in a Sun RPC scenario. [PR1010499](#)

Chassis Cluster

- On all high-end Series devices in a chassis cluster, the security zone is not populated properly on the J-Web interface port configuration page. [PR859200](#)
- On all high-end SRX Series devices in a chassis cluster, the SPC CPU loading of the new backup node might be higher after a data plane failover due to packets forwarded between node0 and node1, which is a dead loop. [PR963033](#)
- On all high-end SRX Series devices in a chassis cluster, the backup node should not send SNMP traps. [PR982777](#)
- On all high-end SRX Series devices in a chassis cluster, when you terminate the GRE tunnel over IPsec VPN, sessions through the GRE tunnel are deleted unexpectedly when the session that is installed on the backup node times out, which is normally at eight times the session timeout. [PR982880](#)
- On all SRX Series devices in a chassis cluster, when source NAT is configured with a port no-translation pool and a port overflow pool with address persistent feature, the port resource of the overflow NAT pool leaks on backup node when the translated IP address creates conflict on the port no-translation pool. [PR991649](#)
- On SRX3400 and SRX3600 devices in a chassis cluster, the FPC0 Minor Errors alarm is raised because of the excessive invalid pkt type errors reported by the Network Processing Card. [PR1008968](#)

Command Line Interface (CLI)

- On all high-end SRX Series devices, in a rare condition, a corrupted memory might be created in data-plane, and then executing the command `show xlr pkt_mbuf`, which is a part of the `request support information` command, results in the flowd process crash. [PR1005067](#)
- On all high-end SRX Series devices, when you configure multiple stream mode under the [security log] hierarchy and one of the stream modes is set to severity warning, the system log traffic on the other streams is stopped. [PR1009428](#)

- On all high-end SRX Series devices, system commit synchronize is not supported. Hence, when you configure it will not be committed due to a configuration lock. [PR1012692](#)
- On all high-end SRX Series devices, the CLI auto-complete does not work for any key words after **set system login class <name> permissions** command. [PR1032498](#)

Flow and Processing

- On all high-end SRX Series devices, under certain conditions, the creation of a multicast leaf session might result in an invalid multicast next hop, which crashes the flow module. [PR921438](#)
- On all high-end SRX Series devices, for IDP, AppSecure, ALG, GTP, or SCTP, the flow serialization impacts session performance. This flow serialization continues even after Layer 7 processing is completed. [PR986326](#)
- On all high-end SRX Series devices, the logical tunnel interface encapsulated Frame Relay is not supported. When you configure logical tunnel interface encapsulated Frame Relay, the flowd process crashes. [PR996072](#)
- On SRX1400 devices, datapath debugging does not capture the system-generated packets. [PR1004074](#)
- On SRX5400, SRX5600, and SRX5800 devices with an SRX5K IOC II, the SRX5K IOC II might send packets out of order, causing end-to-end performance degradation. [PR1007455](#)
- On all high-end SRX Series devices (except SRX1400), fragmented IPsec packets might be out of order after decryption, causing TCP packet retransmission and performance degradation. [PR1013223](#)
- On SRX1400 devices, in a rare condition, SPUs might run into dead loop situation. High CPU usage on SPUs will be seen, and the flowd process will crash in the end. [PR1017665](#)
- On all high-end SRX Series devices, when the central point runs in combo mode on an SPC I card, and when enable-utm-memory and in-line-tap IDP mode are enabled concurrently, the flowd process crashes continuously. [PR1019568](#)
- On all high-end SRX Series devices, in some scenarios, the flowd process might generate core files due to stack overflow while running a log collection script on the device. [PR1020739](#)
- On all high-end SRX Series devices, when heavy load is on the Packet Forwarding Engine management CPU, the speed of IPC sending and receiving between the Routing Engine and the Packet Forwarding Engine might not match, causing security policies to become out-of-sync between RE and PFE. [PR1022351](#)
- On SRX5400, SRX5600, and SRX5800 devices with an SRX5K IOC II, configuring a sampling feature (flow monitoring) might cause high kernel heap memory usage. [PR1033359](#)

General Packet Radio Service (GPRS)

- On all high-end SRX Series devices with GTP enabled, some GTP traffic might be dropped due to the reason message Reason zero TID/TEID. This is because some GTP messages do not contain a TEID value in the GTP message header (such as Identification Response messages), and these messages are dropped incorrectly. [PR999468](#)

Infrastructure

- On SRX1400, SRX3400, and SRX3600 devices configured with firewall simple filters, if you change the simple filter terms, some terms might not be installed properly in the data plane. As a result, the simple filter might not work as expected. [PR1012606](#)

Interfaces and Routing

- On all high-end SRX Series devices, CoS buffer sizes are not recalculated after you delete the interface units. This might result in suboptimal CoS behavior. [PR953924](#)
- On SRX5400, SRX5600, and SRX5800 devices configured with SPC II cards, memory leak might occur on the SPC II Control Plane Processor (CPP), causing the SPC II CPP to reboot. [PR975345](#)
- On SRX5400, SRX5600 and SRX5800 devices, the egress packets delay bandwidth in queue 4 to queue 7 might be dropped when traffic bursts. [PR1007778](#)
- On all high-end SRX Series devices, when a new user is created, the home directory for the user is not created. [PR1015156](#)
- On all high-end SRX Series devices, during route deletion on Packet Forwarding Engine, next-hop entries might not be deleted, these stale next-hops may continue to be used by sessions resulting in flowd process crash. [PR1017037](#)

Intrusion Detection and Prevention (IDP)

- On all high-end SRX Series devices, due to a software defect in XML parsing failure cases, the idpd process might crash during the updating of IDP security packages. [PR1011610](#)

J-Web

- On all high-end SRX Series devices in a chassis cluster, when the switch to Layer 2 mode button is pressed in J-Web, it does not ask for any confirmation and converts to transparent mode immediately and reboots the device. [PR1007740](#)
- On all high-end SRX Series devices, on the Dashboard page, the serial number and the system uptime are not displayed. [PR1009371](#)

Platform and Infrastructure

- On all high-end SRX Series devices, there is some buffer leak in Application Delivery Controller (ADC) and Transparent Load Balancer (TLB) services due to the malfunction of atomic functions. [PR934768](#)

Security

- On all high-end SRX Series devices, when you swap the sequence of security policies or when security policies are disabled by scheduler, the applications configured in these security policies might be added to other enabled security policies, causing unexpected applications to be evaluated by other security policies, and traffic to be permitted or denied unexpectedly. [PR1033275](#)
- OpenSSL released a Security Advisory that included CVE-2014-3566 known as the "POODLE" vulnerability. The SSL protocol 3.0 (SSLv3) uses nondeterministic CBC padding, which makes it easier for man-in-the-middle attackers to obtain clear text data through a padding oracle attack. OpenSSL is upgraded to support for SSL 3.0 fallback protection (TLS_FALLBACK_SCSV). Refer to JSA10656 for more information. [PR1033938](#)

System Logging

- On all high-end SRX Series devices, the custom dynamic group with the service TCP filter or UDP filter does not include TCP or UDP port-bound attack signatures. The following error message is displayed:

'dynamic-attack-group OTHER-PROTO-REC-CTS'

Attack TCP-PROTO-REC-CTS: No matching members found. Group is empty error: configuration check-out failed

The group should not be empty, because of the configured queries of the custom dynamic group. [PR1002526](#)

- On all high-end SRX Series devices, RT_PFE errors might be generated due to reroute failure when a more specific route entry is added or deleted. [PR1009947](#)
- On all high-end SRX Series devices, the **flowd_octeon_hm: pconn_client_connect: Failed to connect to the server after 0 retries** message repeats in the log. [PR1035936](#)

Unified Threat Management (UTM)

- On all high-end SRX Series devices, due to a memory leak issue in the utmd process, the utmd process might cause control plane CPU utilization that is higher than expected even when the Unified Threat Management (UTM) feature is not enabled. The memory leak can only be triggered if there is a UTM license installed on the system. [PR1027986](#)
- On all high-end SRX series devices with Web Trends Enhanced Log File (WELF) format configured for security log, when system generates very long WELF formatted logs (such as, more than 1k bytes), and it is truncated on Packet forwarding engine (PFE) and sent to Routing Engine (RE), a memory corruption issue might occur in this situation, resulting in the flowd process to crash. This issue occurs more when Unified Threat Management (UTM) Web Filtering is configured. [PR1038319](#)

VPN

- On all high-end SRX Series devices, the IPv6 traffic is reordered during the encryption of IPsec VPN because the fragment order is not maintained for the IPv6 traffic. [PR962600](#)
- On all high-end SRX Series devices, dynamic VPN user groups are not able to access certain remote resources. In this scenario, there are two policies referring to the same dynamic VPN and one of the policy directions is not set. Hence, the lookup fails in the null policy at the end. [PR988263](#)
- On all high-end SRX Series devices with IPsec VPN configured, due to a rare timing issue, IPsec VPN traffic might be dropped as the reason of "bad SPI" on the traffic receiving side during IPsec Security Association (SA) rekey. [PR1031890](#)
- On all high-end SRX Series devices with policy-based IPsec VPN configured, deleting security policies that are associated with a VPN tunnel might result in a stale VPN tunnel remaining. In addition, the stale VPN tunnel might be associated with the newly added security policies. [PR1034049](#)

Resolved Issues - 12.1X46-D25***Application Layer Gateways (ALGs)***

- On all high-end SRX Series devices, when RTSP ALG traffic passes through the routing instance type virtual router, under some conditions the traffic is dropped. [PR979899](#)
- On all SRX Series devices, when there is heavy SIP traffic through the device, high CPU usage is seen on one or more SPUs. This issue occurs due to a certain type of SIP-handling logic, which dumps payload packets to the internal buffer. This logic has been optimized to reduce load on the SPU. [PR985932](#)

Chassis Cluster

- On all high-end SRX Series devices in a chassis cluster with multicast configuration, when the Redundancy Group 0 (RG0, a Redundancy Group for RE) failover, it might cause too many memory fragments in kernel, and result in some control operation failure due to lack of continuous memory. [PR944604](#)
- On all high-end SRX Series devices in a chassis cluster, when the secondary node becomes ineligible due to control link failure it might still forward the traffic. This causes the reth interface to flap and the related traffic to drop when the secondary node is in ineligible state. [PR959280](#)
- On all SRX Series devices in a chassis cluster with the PPTP ALG enabled and the PPTP session closed, a memory corruption might occur on the secondary node, which causes the flowd process to crash. [PR993447](#)
- On all high-end SRX Series devices in a chassis cluster with interface monitoring enabled, interfaces might be incorrectly monitored as down due to a memory allocation issue. [PR1006371](#)

Command-Line Interface (CLI)

- On all high-end SRX Series devices, the **show interface extensive** command is cut short with the error message **error: route rpf stats get for interface**. [PR930630](#)

Dynamic Host Configuration Protocol (DHCP)

- On all high-end SRX Series devices, you cannot get the DHCP relay information through SNMP if DHCP relay is configured under the logical system. For example: **bash-3.2# snmpwalk -c lsys1/default@junos -t 5 -v 1 -Os -Oq -Oe -Pu -m /tmp/jnx-smi.mib:/tmp/jnx-jdhcp.mib 10.208.131.136 jnxJdhcpRelayStatistics bash-3.2#**
[PR909906](#)
- On all high-end SRX Series devices, DHCPv6 does not work in IPv6 mode. [PR942246](#)
- On all high-end SRX Series devices which work as a DHCP server, if the server receives a DHCP INFORM packet from a binding client, and then this binding entry is released by issuing **clear system services dhcp binding** command, or the server receives a DHCP RELEASE packet from the same client, this will cause the IP address not get released and the same IP address might assign to a different client in the subsequent assignment. [PR969929](#)

Flow and Processing

- On all SRX Series devices, when you run the **clear security flow session** command with a prefix or port filter, some of the sessions are not matched with the filter, causing a traffic drop or delay. This issue is triggered by any of the filters. [PR925369](#)
- On all high-end SRX Series devices, in certain situations, flow sessions time out and get corrupted. This leads to the flow sessions being set to an abnormally high value, which eventually leads to the session table becoming full. [PR955630](#)
- On all high-end SRX Series devices, when you configure an ICMP **probe-server** option under the [**services rpm**] hierarchy for a specific interface (for example, ge-0/0/0), the device does not respond to ICMP requests from this interface. Other interfaces are not affected and continue to respond to ICMP requests. [PR960932](#)
- On all high-end SRX Series devices, when you reboot the passive node, the CPU usage increases on flow SPUs of the primary node and this lasts for a few seconds when the traffic latency is increased. [PR962401](#)
- On all high-end SRX Series devices, filter-based forwarding (FBF) rules are ignored when existing sessions are rerouted. [PR962765](#)
- On all high-end SRX Series devices deployed in a multicast scenario, a memory leak on the fwdd process might occur when the multicast routes change. [PR963116](#)
- On all high end SRX devices, when it processes fragmented packets, the first fragment (the fragment contains layer 4 information) will be used to create session, and the subsequent fragments will be queued on a memory block. When in session creation stage, the queued fragments might be processed for flow processing even though the session is still in pending state, this results in the order information lost and the fragmented packets forwarded out of order. [PR993925](#)

Hardware

- On SRX5400, SRX5600, and SRX5800 devices configured with SPC II cards, memory leak might occur on the SPC II Control Plane Processor (CPP), causing the SPC II CPP to reboot. [PR975345](#)
- On SRX5400, SRX5600, and SRX5800 Series devices with SPC used, in certain condition, SPUs might hang due to memory unaligned accessing. Memory unaligned accesses is supported by default. [PR980122](#)
- On SRX5400, SRX5600, and SRX5800 devices, after fabric reconnect (it can be reconnected by issue the **restart chassis-control immediately** command), setting the fabric plane to offline and then setting it to online will fail. The fabric plane link error message will be seen by issue the **show chassis fabric fp** command. [PR990679](#)

On all high-end SRX devices, session ager might gets stuck due to a memory corruption, causing maximum session limitation to be reached on services processing units (SPUs). [PR991011](#)

Interfaces and Routing

- On SRX5400, SRX5600, and SRX5800 devices, there are incorrect counters on reth interface. [PR978421](#)

Intrusion Detection and Prevention (IDP)

- On all high-end SRX Series devices, when the LACP mode is fast and the IDP is in inline-tap mode, a LACP flap might occur when you commit the configuration. [PR960487](#)
- On all high-end SRX Series devices, when the IDP security package update contains a detector version change, the configured detector kconst values are not pushed from the idpd process to the Packet Forwarding Engine. Hence, the newly loaded detector takes default values. [PR971010](#)
- On all SRX Series devices, when you configure an automatic security package update without configuring the schedule interval and start time, high CPU usage on the idpd process is seen. [PR973758](#)

Network Address Translation (NAT)

- On all high-end SRX Series devices, in rare cases, the device starts using sequential source ports for source NAT because of random function memory corruption. [PR982931](#)

Screens

- On all high-end SRX Series devices with flooding type screens configured, if multiple logical interfaces on the same Network process Unit (NPU) have been configured in the same zone, then changing the flooding thresholds might cause each of these logical interfaces to have inconsistent thresholds, and sometimes some logical interfaces might not have any screen flood protection at all. [PR972812](#)

System Log

- On all high-end SRX Series devices, every time a user logs in with SSH, a **verixec: fingerprint mismatch** message is reported in the log. [PR929612](#)
- On all high-end SRX Series devices, the new entry or flag representing an alert notification is seen in the system log message. If the alert is configured in the IDP rules, the flag is set to yes; otherwise, it is set to no. [PR948401](#)
- On all high-end SRX Series devices, **Duplicate FLOW_IP_ACTION** logs are generated while sending traffic. [PR959512](#)
- On all high-end SRX Series devices, the SNMP walk for the jnxPicType2ASPCXLP object might fail and show the jnxPicType2ASPCXLP (could not resolve 'jnxPicType2ASPCXLP' to an OID) error message in the logs, and fails to receive information from the device. [PR974463](#)

Virtual Private Networks (VPNs)

- On all high-end SRX Series devices, in certain situations when the device has more than one IKE Security Association (SA) installed for the same peer device and DPD is triggered, the messages are not sent out from the device to the peer device, causing the IKE SA to be installed on the device until the IKE SA expires. [PR967769](#)
- On all high-end SRX Series devices, when the device is configured with similarly named CA profiles (example: caprofile, caprofile_1, caprofile_3 and so on) and CA certificates are loaded to these profiles, when first CA certificate is cleared other certificates which has the CA profile that starts with the same keyword will be cleared as well. [PR975125](#)

Resolved Issues - 12.1X46-D20

Application Layer Gateways (ALGs)

- On all high-end SRX Series devices, the Microsoft Active directory or Microsoft Outlook client might get disconnected from the server because the MS-RPC ALG incorrectly drops the data connections under heavy load. [PR958625](#)

AppSecure

- On all high-end SRX Series devices, the application firewall module might cause the Network Security Daemon (NSD) to create up to 4 KB of memory leak when you commit each configuration. [PR969107](#)

Dynamic Host Configuration Protocol (DHCP)

- On all high-end SRX Series devices, DHCPv6 does not work in the IPv6 mode. [PR942246](#)

Flow and Processing

- On all high-end SRX Series devices, the flowd process might crash during the session installation. [PR956775](#)

J-Web

- On all high-end SRX Series devices, J-Web does not accept the keyword “any” in the address-book object name. [PR944952](#)

Network Address Translation (NAT)

- In Junos OS Release 12.1X46-D10 and earlier, the device could not send the SNMP trap for the NAT pool with logical systems configured. Starting in Junos OS Release 12.1X46-D20, the SNMP trap for the NAT pool with logical systems configuration can be sent from the device. [PR959219](#)

Platform and Infrastructure

- On all high-end SRX Series devices, if the NTP server is not a stratum 1 server, the NTP synchronization process cannot be completed. To confirm this issue is occurring, use the **show ntp status** command. [PR864223](#)
- On all high-end SRX Series devices, the nsd process might hold a buffer related to the NAT proxy-arp process, and it does not release the buffer. This causes a memory leak on the nsd process when you commit a configuration. [PR931329](#)
- On all high-end SRX Series devices, in certain circumstances, the high CPU consumption on the data plane and an eventual exhaustion of the internal system buffers might corrupt the forwarding table, which causes the traffic to drop partially. [PR938742](#)
- On SRX5600 and SRX5800 devices, during the LICU code upgrade for the control port, the FPCx (DPC) changes to any erroneous number and needs to use the non-IOC port (SPC, existing or not) on the chassis.

Refer to KB17947 for additional information. [PR953029](#)

System Log

- On all high-end SRX Series devices, the error **OpenSSL: error:14090086:lib(20):func(144):reason(134)** means that server certificate verification has failed. The certificate might be a self-signed certificate or an expired certificate. [PR932274](#)

Unified Threat Management (UTM)

- On all high-end SRX Series devices, when you install a license, you might see the message **license not valid for this product add license failed**. Even though the message appears, the feature still functions normally. In addition, the **show system license** command does not display the Sophos antivirus, antispam, or Web filtering licenses. [PR948347](#)
- On all high-end SRX Series devices, UTM blacklists and whitelists should work without an EWF license. [PR970597](#)

Virtual Private Networks (VPNs)

- On all high-end SRX Series devices, during VPN configuration change with an interface configuration change at the same commit, or after rebooting the device with VPN and interface configured together, the tunnel sessions created in flowd are missing. This impacts the traffic flow on that tunnel. The invalid bind interface counter returns a nonzero value when you run the **show usp ipsec global-stat** command. [PR928945](#)
- On SRX3400, SRX3600, SRX5400, SRX5600, and SRX5800 devices, high CPU usage occurs after installing additional SPC cards without a full cluster reboot and IPsec tunnels carry the SCTP traffic anchored on the device. [PR945162](#)
- On all SRX Series devices, any configuration changes to the st0.x interface might delete the NHTB entries for unrelated st0 interfaces. [PR958190](#)

Resolved Issues - 12.1X46-D15

Application Layer Gateways (ALGs)

- On SRX Series devices with the VoIP-related ALG (either H.323 or SIP) and NAT enabled for the VoIP traffic, the corresponding ALG creates persistent-nat-binding entries for the reverse VoIP traffic (even though the persistent NAT feature is not configured in the source NAT rule) when VoIP traffic is transmitted into a custom routing instance. Hence, the system does not apply the custom routing instance information to the persistent-nat-binding entries, and the reverse traffic that matches the persistent-nat-binding entries is forwarded to the default routing instance instead of to the custom routing instance. The reverse traffic is dropped or forwarded to the wrong place. [PR924553](#)

Chassis Cluster

- On devices in a chassis cluster working as a Unified Access Control (UAC) enforcer, when RGO failover occurs, the Packet Forwarding Engine might connect to the uac process before the uac process connects to the UAC server. In this condition, the uac process conveys to the Packet Forwarding Engine that the UAC server is disconnected. When the Packet Forwarding Engine receives this information, it denies new traffic that matches the UAC policies. The traffic is resumed after the connection of the uac process and UAC server is established. [PR946655](#)

Dynamic Host Configuration Protocol (DHCP)

- On all high-end SRX Series devices, after you configure DHCPv6 in IPv6 mode, the dhcpv6 process crashes. [PR940078](#)

Flow and Processing

- For SCTP IPv6 traffic in traffic logs, all the source and destination ports are marked as using port 1. [PR928916](#)
- When IKE packets are received before Junos OS default applications are pushed to the Packet Forwarding Engine, the IKE sessions will be established without the IKE application having been marked. As a result, the fragmented IKE packet cannot be sent to iked, because the IKE session has not used IKE applications. [PR942730](#)
- On devices with 1 GB of memory, if the advanced services license is configured with the **reduce-dp-memory** option, memory is not released from the data plane to the control plane. [PR895648](#)

Interfaces and Routing

- When IS-IS is configured between the SRX device and some third-party devices, after the SRX device is rebooted and the IS-IS adjacency is reestablished, the routes advertised by the third-party devices might not install into the routing table in some cases. [PR935109](#)

Intrusion Detection and Prevention (IDP)

- On SRX Series devices configured with IDP, for the AppSecure, ALG, GTP, or SCTP features that require the serialization flow processing, the memory buffer might leak, causing the flowd process to crash. [PR930728](#)

J-Web

- J-Web does not accept the address if the object name includes the word “any”. [PR944952](#)

Network Address Translation (NAT)

- In Junos OS Release 12.1X46-D10 and earlier, the device could not send the SNMP trap for the NAT pool with logical systems configured. Starting in Junos OS Release 12.1X46-D15, the SNMP trap for the NAT pool with logical systems configuration can be sent from the device. [PR959219](#)

System Log

- An illegal pointer address generates eventd core files. [PR784037](#)

Unified Threat Management (UTM)

- EWF logs are not marked with user role information. [PR936799](#)

Virtual Private Networks (VPNs)

- On all SRX Series devices configured with IPsec VPN and the VPN monitoring is enabled, the VPN monitoring function triggers socket leak, and it might result in some critical issue, such as SPUs unresponsive. [PR940093](#)
- On all SRX Series devices, when IPsec is used in a chassis cluster, after the SPU or flowd uptime reaches 50 days or more, the amount of RTO traffic on the fabric link increases. [PR941999](#)
- On all SRX Series devices with multiple proxy-identity (MPID), the dead routes are seen while moving the st0 interface from one virtual router to another. [PR943577](#)
- After traffic-selector configuration is deleted from the VPN configuration object, the data traffic stops passing through the tunnel. [PR944598](#)
- SRX Series devices cannot proceed to automatic certificate reenrollment through SCEP. The certificate validity period is incorrectly calculated during the autorenewal

process. Also, when the CRL is downloaded through LDAP, it can be partially received from the CA server and the pkid process goes up. [PR946619](#)

- When there are more than 100 traffic selectors configured on a VPN configuration object along with configured, established, tunnels, if all IPsec SAs for this VPN configuration object are cleared at the same time (because of a configuration change on a peer or the use of the clear operational command), the bind-interface associated with that VPN configuration object might be marked as down. [PR947103](#)

Resolved Issues - 12.1X46-D10

Application Layer Gateways (ALGs)

- The b attribute (pertaining to bandwidth) in a SIP Session Description Protocol (SDP) message was not carried forward after SIP ALG processed the packet. [PR875211](#)
- When an RTSP TCP segment cannot be processed because it is too small or incomplete, the RTSP ALG holds it and waits for the next segment. An RTSP endpoint does not receive an ACK for segments that are too small, so it retransmits the segment several times. Eventually, the RTSP endpoint resets the TCP connection. [PR887601](#)
- With RTSP ALG traceoption enabled, if failover occurs on the device, it will trigger flowd crash that RTSP ALG receiving interleaved RTSP traffic before the RTSP objects are synchronized completely. [PR893136](#)
- In rare cases when ALG is used for flow processing MSS (Maximum Segment Size) in TCP 3-way, handshake is announced in one direction with value higher than 32,120, the next packets in opposite direction get window size value reduced to 0. [PR895498](#)
- In certain circumstances, if the OPTIONS method is used to create a call, and the INVITE method is used to reuse the call, the SIP ALG would apply an incorrect state. As a result, the device might drop the ACK of 200-OK. [PR898956](#)
- The SCTP module drops the SCCP packet when the received SCCP pointer goes out of order. [PR901584](#)
- On devices enabled with the MS-RPC ALG, the flowd process might crash frequently when heavy MS RPC traffic is processed by the MS-RPC ALG. [PR907288](#)

AppSecure

- AppID is using order to selectively report nested applications that it has matched in different transactions on the same session. This means that only nested applications with a higher order are reported. The expected behavior is that it should report nested applications as and when it detects them in the transaction. [PR914567](#)

BGP

- Under specific time-sensitive circumstances, if BGP determines that an UPDATE is too big to be sent to a peer, and immediately attempts to send a withdraw message, the routing daemon (rpd) may crash. An example of an oversized BGP UPDATE is one where a very long AS_PATH would cause the packet to exceed the maximum BGP message size (4096 bytes). The use of a very large number of BGP Communities can also be used to exceed the maximum BGP message size.

Please refer to JSA10609 for additional information. [PR918734](#)

Certificate Authority (CA) Profile

- When you run the **show security pki *-certificate** command, the result displays time without a time zone. [PR746785](#)

Chassis Cluster

- On devices in a chassis cluster with the second control link connected, when CRM is installed, and the primary node is power-cycled, the primary node takes over RG-0 ownership when the primary node is rebooted. [PR679634](#)
- On devices in a chassis cluster, if a reth Layer 3 logical interface is disabled, the reth interface remains active and the direct route for this logical interface is not removed from the forwarding table. All the traffic destined for the disabled network still gets routed out to the disabled reth interface. [PR740856](#)
- On devices in a chassis cluster, when you execute the **clear system commit** command, it clears commit only from the local node. [PR821957](#)
- On devices in a chassis cluster, during a control link failure, if the secondary node is rebooted by control link failure recovery, the rebooted node goes into disabled state even after startup. [PR828558](#)
- On SRX1400, SRX3400, and SRX3600 Series devices, under certain conditions, the em0 (tsec1) detection and recovery mechanism is not working as expected. This might cause the chassis cluster to fail, a “split-brain” condition to occur, or all FPCs to be reset on the local node.



NOTE: Do not use the security policy count and make sure trace options are disabled. Do not use **set security log mode event** command; instead use **mode stream** (default mode).

[PR877604](#)

- On devices in a chassis cluster, the chassisd log outputs are flooded with the following message: **LCC: fru_is_present: out of range slot -1 for SCB**. [PR889776](#)
- On devices in a chassis cluster, in certain IPv6 configurations, the SPU sends out packets with an invalid header on the secondary node, which in turn triggers the hardware monitoring failure on the secondary node. [PR935874](#)

Command-Line Interface (CLI)

- There is no specific CLI command to display the count of sessions allowed, denied, or terminated because of UAC enforcement. [PR733995](#)
- The **show security pki *-certificate** shows the time without a time zone. [PR746785](#)
- The output of the **show security pki ca-certificate detail** command includes the Auto-re-enrollment section. This is incorrect because automatic reenrollment is not supported for CA certificates. [PR877574](#)
- Certain combinations of Junos OS CLI commands and arguments have been found to be exploitable in a way that can allow root access to the operating system. This may allow any user with permissions to run these CLI commands the ability to achieve elevated privileges and gain complete control of the device.

Please refer to JSA10608 for additional information. [PR912707](#), [PR913328](#), [PR913449](#), [PR913831](#), [PR915313](#), [PR915957](#), [PR915961](#), [PR921219](#), [PR921499](#)

Flow and Processing

- When DNS ALG was enabled, the rewrite rules applied on the egress interface might not work for DNS messages. [PR785099](#)
- On all high-end SRX Series devices, when plugins that use TCP proxy (such as ALGs or UTM) are configured, a certain sequence of valid TCP packets crashed the flow daemon (flowd). Repeated crashes of flowd represented an extended denial of service condition for the device. [PR791201](#)
- On all high-end SRX Series devices, when fragmented jumbo frames are reassembled in the SPU (reassembling might be required by IDP feature, ALG feature, ESP/AH packets, and L2TP packets) and if the size of the reassembled packet becomes larger than 9712 bytes, the packet is dropped in the internal device, and the device reports XLR egress packets corruption issues. [PR819621](#)
- On all high-end SRX Series devices, the SPU level kernel crashed and generated vmcore files when processing traffic that required serialized packet processing in some application modules such as IDP, ALGs, application security, and so on. [PR855397](#)
- Current implementation of timeout for http is 1800s, the default timeout should be 300s. [PR858621](#)
- Periodic multicast packets such as NTP do not refresh the route, and packets are dropped intermittently. [PR869291](#)
- On SRX Series devices, during ARP floods of the data plane Packet Forwarding Engine, the CPU spikes might impact transit and host-bound traffic. [PR871704](#)

- On devices in a chassis cluster, after data plane RG1 failover, the RTSP data packet is queued, and a duplicate RTSP data packet is processed by the device; the flowd process crashes and generates core files. [PR883397](#)
- When TCP SYN flood protection is enabled and triggered, and if the Window Scaling option is used between a TCP client and server, TCP communication is reset abnormally. [PR886204](#)
- On all high-end SRX Series devices, due to incorrect computation of central point IPv6 sessions, the output of the total central point sessions is incorrect for the **show security monitoring fpc number** command. This is only a display issue and does not affect actual central point sessions or the traffic passing through. [PR888890](#)
- On SRX1400 devices, the egress packets are dropped. There is an increase in the number of egress packets dropped when the traffic passes through the ports of the SRX1K-SYSIO card. [PR899184](#)
- When flow traceoptions are used to debug source NAT traffic, packet filter did not work. This resulted in a large amount of unexpected traces. [PR905568](#)
- The CRL download fails for fragmented LDAP packets. [PR910947](#)
- On all high-end SRX Series devices, when you delete a large number of interfaces and commit, and immediately add a large number of interfaces and commit, the session scan might fail. The session related to the deleted interface might still be active, in which case the subsequent traffic drops if it matches the old session. This occurs in a scenario when the deleted interface is added back on the “immediately add” action, and the remote host still generates the traffic matching the session. This issue occurs as the session interface is detected in invalid state in flow checking. [PR915422](#)
- J-Flow might not work as expected; the cflowd packets are not seen for version 5 and version 8 sampled flows. [PR916986](#)

General Packet Radio Service (GPRS)

- If a GTPv1 user plane (GTPv1-U) tunnel update conflicts with a secondary tunnel, then core files are generated. [PR888067](#)
- When there is inconsistency in the NAT rule configuration for the IP address in the IP header and in the GTP payload, packets are dropped.
 - When there is a NAT rule for the IP address in the GTP payload and no NAT rule for the IP address in the IP header, the tunnel is set up on a wrong SPU, and the control and data traffic on the tunnel might be dropped.
 - When there is a NAT rule for the IP address in the IP header and no NAT rule for the IP address in the GTP payload, the packet is dropped to keep the consistency of the NAT rule configuration.

[PR921313](#)

Hardware

- When the device is rebooted, the next-generation SPC card might not boot up due to I2C bus hang. Error messages related to "I2C" errors also appear in the log. [PR923255](#)

Infrastructure

- On all high-end SRX Series devices, when the device authentication is through RADIUS server and the password protocol is Microsoft CHAP version 2, the password change operation fails as the user password change is enforced through Microsoft Active Directory server. [PR740869](#)
- After an upgrade, you cannot copy files between nodes in a cluster using the **file copy** command. [PR817228](#)
- In a DHCP-relay subscriber management environment, with an output firewall filter configured on an IRB interface to discard the DHCP offer packets, while DHCP-relay subscribers log in, the Junos OS kernel tries to free an already freed memory buffer, which causes the kernel to crash and generate core files. [PR824470](#)
- When the backup Routing Engine kernel fails, some devices send a message to the master Routing Engine to generate a core file. [PR854501](#)
- On SRX1400 devices with 10-Gigabit Ethernet, when the system I/O card is inserted on SFP-T of ge-0/0/7, ge-0/0/8, or ge-0/0/9 interface, the device interface LEDs light immediately. [PR865899](#)
- If the secondary control link (em1) interface uses SFP-T, the interface is down when you add node 1 to the cluster. [PR873253](#)
- On devices in a chassis cluster, after control plane Redundancy Group (RG0) failover, SFPs might have more if states than the new master Routing Engine. This difference leads to sequence number mismatch and causes cold synchronization failure, and all FPCs might reboot. After the FPCs reboot, a "split brain" situation occurs in which both nodes become primary. [PR885889](#)
- E2edebg traces are not generated for all the events. [PR919471](#)

Interfaces and Routing

- On the K2-Routing Engine (64-bit Routing Engine) when speed or link mode are statically configured on the device for the fxp0 interface, the driver for fxp0 accepts the configuration from DCD process. The K2-Routing Engine does not propagate the setting to the hardware driver. Instead, the driver setting is forced to auto-negotiate. Thus, as the fxp0 interface is auto-negotiating, and the far end device is forced to 100/full, the auto-negotiation on fxp0 will detect the speed but will not detect the duplex and hence, defaults that duplex to half-duplex. [PR704740](#)
- On VLAN tagged Ethernet frames (802.1p), you cannot modify the VDSL priority bits. [PR817939](#)
- Multicast stream is redirected to other member links on the ae interface or on the reth LAG even when the link in use is disabled. [PR867529](#)

- When a SHDSL Mini-PIM is configured in 2-wire mode with annex mode as Annex B/G, one of the physical interfaces does not come up. [PR874249](#), [PR882035](#)
- On devices in a chassis cluster, when a session created as the incoming interface is a VPN secure tunnel interface (ST interface) and the outgoing interface is a logical tunnel interface (LT interface), this session is incorrectly marked as active on the secondary node. When this session expires on the secondary node, the sessions on both cluster nodes might get deleted and interrupt the traffic. [PR896299](#)
- When multiple routing instances are defined, DNS names in the address-book entries might not get resolved. This results in corresponding security policies to be nonoperational. [PR919810](#)
- When multiple IP addresses from an overlapping subnet are configured on a single interface, the interface enable-related or disable-related changes might not work. [PR920993](#)

Intrusion Detection and Prevention (IDP)

- On XLP platforms, setting the **max-sessions** option in an application identification configuration did not impact the attack traffic. [PR809384](#)
- After the Junos image is upgraded, we recommend that you download a completely updated IDP security package and then perform the installation. Subsequent incremental updates (default) work fine. If a complete update is not performed, the device might end up adding only the new signatures downloaded in incremental order, leaving the device unprotected from a large set of signatures. [PR876764](#)
- On SRX Series devices with IDP enabled, if IDP exempt rule is configured, a change of the IDP rule configuration (such as a change to source or destination, action, or signature) might cause the flowd process to crash and core files are generated. [PR877865](#)
- When there are a large number of ASC entries (100,000 or more), and the entries are listed using CLI command, the flowd process might crash. [PR886173](#)
- On all high-end SRX Series devices, maximize sessions inline-tap equal mode is not supported in Junos OS Release 12.1X46-D10. If the maximize sessions inline-tap equal mode is configured in a release earlier than Junos OS Release 12.1X46-D10, when you upgrade to Junos OS Release 12.1X46-D10, the configuration changes to maximize sessions inline-tap firewall mode. [PR889597](#)
- On SRX Series devices, the flowd process might crash when IDP is enabled using software based pattern matching and detects more than one attack entry for the same attack. [PR907703](#)

J-Web

- The J-Web interface was vulnerable to HTML cross-site scripting attacks, also called XST or cross-site tracing. [PR752398](#)
- In J-Web, SRX Series devices fail to downgrade from Junos OS Release 12.1X46-D10 through HTTP file upload. [PR918112](#)

- In J-Web, if the policy name was "O", the penultimate-hop popping (PHP) function treated it as empty, and traffic log output could not be viewed. [PR853093](#)
- In J-Web, the LSYS operation might cause MGD to generate a core file, and **compare before commit** does not work. [PR889029](#)

Network Address Translation (NAT)

- Under certain conditions, a duplicate SNMP index might be assigned to different interfaces by the kernel to the mib2d (Management Information Base II daemon). This might cause mib2d and other processes such as lacpd (LACP daemon) to crash and generate core files. [PR836823](#)
- On devices enabled with the PIM protocol, the flowd process crashed and generated core files, when there was a unicast PIM register message received with encapsulated multicast data; and if NAT process was involved in the session for the received PIM packet. This issue was observed on standalone high-end SRX Series devices, and on devices in a chassis cluster. In the case of devices in a chassis cluster, the flowd process crashed on both node 0 and node 1. [PR842253](#)
- In a root system, the destination and static NAT rule cannot send system log and trap messages when the number of sessions reaches the threshold value. In a logical-system, the source, destination, and static NAT rule cannot send system log and trap messages when the number of sessions reaches the threshold value. [PR905359](#)
- On devices in a chassis cluster, the chassis cluster rule number of sessions in the SNMP query or walk result is the sum of the real number of sessions of the primary node and the secondary node. [PR908206](#)
- On all high-end SRX Series devices, when source NAT is configured with persistent NAT enabled, sometimes the persistent NAT bindings leak on the central point. [PR910116](#)

Routing Policy and Firewall Filters

- If more than 10 virtual routers (routing instances) or logical systems (LSYS) are configured on a device, DNS fails to resolve addresses. A maximum of only 10 routing instances and LSYS can be configured per DNS name server. [PR896174](#)

Screen

- On all high-end SRX Series devices with IP spoofing screen enabled, the routing table search fails when it is locked by the system. As a result, false positives occur on IP spoofing detection. [PR901507](#)
- On all high-end SRX Series devices, security screen are not allocated for more than 165 zones due to memory limitation. If a security screen is enabled for more than 165 zones, only 165 zones are actually enabled and the memory is exhausted by the screen allocation, resulting in traffic interruption. [PR913052](#)

Security

- The glob implementation in libc allows authenticated remote users to cause a denial of service (CPU and memory consumption) via crafted glob expressions that do not match any pathnames. This vulnerability can be exploited against a device running

Junos OS with FTP services enabled to launch a high CPU utilization partial denial of service attack.

Please refer to JSA10598 for additional information. [PR558494](#)

- If Proxy ARP is enabled on an unnumbered interface, an attacker can poison the ARP cache and create a bogus forwarding table entry for an IP address, effectively creating a denial of service for that subscriber or interface. When Proxy ARP is enabled on an unnumbered interface, the router will answer any ARP message from any IP address which could lead to exploitable information disclosure.

Please refer to JSA10595 for additional information. [PR842092](#)

System Log

- On SRX3400 and SRX3600 devices, the following system logs are seen in the messages file:

sfchip_show_rates_pfe: Fchip Plane 0, dpc 0, pfe <1/2/3>: Invalid dpc

These system logs do not affect the device. [PR738199](#)

- Fetching ppX interface statistics leaks in pfestat_table leads to the following error logs:

pfestat_req_add: pfestat table out of ids

During this state it is not possible to fetch any interface statistics. [PR751366](#)

- On all high-end SRX Series devices, when a Routing Engine card is removed and placed again, swapped, or rebooted, the following error message appears in the system log for an hour:

No response from the other routing engine for the last 300 seconds

[PR875189](#)

- SRX5600 and SRX5800 devices with an SRX5K-SPC-4-15-320 (next-generation SPC) might generate one of the following system logs on the messages file:

spu_mac_get_linkstate:spu (<fpc#>/<pic#>) – phy link<link#> failed

spu_mac_get_linkstate:%PFE-3:(<fpc#>/<pic#>) –MAC layer link failed

In this condition, the affected SPU cannot do any flow processing until the system is rebooted. [PR914736](#)

- On SRX1400 devices in some cases, the traffic gets interrupted for about 5 seconds occasionally and the following log message appears:

XLR ingress pause

[PR921692](#)

- The session ID of AppTrack logs does not include the SPU ID. Hence, there is a mismatch with firewall log session ID and AppTrack log session ID of the same session. The AppTrack log now has the same session id used in the firewall logs.[PR924941](#)

Unified Threat Management (UTM)

- The enhanced Web filter parser mishandles the URL and host from the HTTP header. This results in an “uncategorized” EWF reply. [PR862602](#)

Virtual Private Networks (VPNs)

- On a SRX Series device, when a session is closed because the user for that session has signed out from the Junos Pulse, the session close log shows the role information as “N/A”. [PR689607](#)
- An IPsec policy for a VPN can contain proposals with different protocol types (ESP or AH). This means that an IPsec SA can be established with either ESP or AH, depending on the protocol type of the peer’s proposal. [PR843281](#)
- When IPsec VPN Internet Key Exchange (IKE) traffic passed through the device, memory leaks were observed and the VPN connection could not be established. [PR857013](#)
- On all high-end SRX Series devices, the Junos Pulse client has been updated from Release 2.0R3 to 4.0R2. [PR868101](#)
- File Descriptor leak occurs during the network-security-trace process when commit configuration changes are made in the **edit security ike** configuration. Eventually, the system reaches the maximum file limit, which results in a system-unmanageable condition. [PR893017](#)
- In a site-to-site IPsec VPN deployments using IKEv2, when tunnels are removed through configuration change, the information is not propagated to the remote peer. Later, when the peer initiates a normal Phase-1 re-key process, the kmd process crashes and core files are generated. [PR898198](#)

Related Documentation

- [New and Changed Features on page 95](#)
- [Known Behavior on page 135](#)
- [Documentation Updates on page 183](#)

Documentation Updates

This section lists the errata and changes in Junos OS Release 12.1X46 documentation.

[Documentation Updates for the Junos OS Software Documentation](#)

This section lists improvements and outstanding issues with the software documentation.

Junos OS for SRX Series Documentation

The Junos OS for SRX Series technical documentation set has been expanded, restructured, and retitled in Junos OS Release 12.1X46-D10 to make it more comprehensive, easy-to-use, and intuitive. Highlights:

- (New) The Complete Software Guide consolidates all of the release-specific content that applies to Junos OS for SRX Series devices (except release notes) into a three volume set of PDFs that you can download and view offline. The first volume contains getting started and administration information; the second contains feature information; the third contains developer information. You can find the PDFs in the Downloads box on the right side of the *Junos OS for SRX Series Services Gateways, Release 12.1X46* index page.
- (New) The *Getting Started Guide for Branch SRX Series* describes how to get up and running with branch SRX Series devices.
- (Expanded) The *Junos OS Monitoring and Troubleshooting Library for Security Devices* contains significantly more content to help network and security managers keep their SRX Series devices running smoothly in their production environments.
- (Expanded) The *Junos OS for SRX Series Services Gateways, Release 12.1X46* index page has been expanded to serve as a “one stop shop” for all of your Junos OS for SRX Series technical documentation needs.

Junos OS Release Notes

In Junos OS 12.1X46-D10 Release Notes and Junos OS 12.X46-D15 Maintenance Release Notes, the SCCP ALG feature description has the following incorrect information:

Support for SCCP v20—This feature is supported on all SRX Series devices.

Starting in Junos OS Release 12.1X46-D10, the SCCP ALG supports version 20. In SCCP v20, several SCCP messages have been updated with a new format.

The correct information is as follows:

Support for SCCP v20—This feature is supported on all SRX Series devices. Starting in Junos OS Release 12.1X46-D10, the SCCP ALG supports SCCP versions 16, 17, and 20 and several SCCP messages have been updated with a new format. Cisco Call Manager (CM) version 7 uses SCCP version 20.

Administration Guide for Security Devices

- The following note is added to the Administration Guide for Security Devices, in the Encrypting Configuration Files topic:



NOTE: The `request system set-encryption-key` command is not supported on high-end SRX devices, therefore, this task does not apply to such devices.

- Under the Configuration tab, the “Minimum DHCP Local Server Configuration” topic has been updated to replace the pool name and group name with more appropriate names. The text should read as follows:

```
[edit access]
address-assignment {
  pool acmenetwork family inet {
    network 192.168.1.0/24;
  }
}

[edit system services]
dhcp-local-server {
  group mobileusers {
    interface ge-1/0/1.0
  }
}

[edit interfaces ge-1/0/1 unit 0]
family {
  inet {
    address 192.168.1.1/24
  }
}
```

Application Identification Feature Guide for Security Devices

- Under the Administration tab, in the example titled “Example: Creating a Configuration Workflow for SSL Proxy,” there is an incorrect **openssl** command. In Step 2d of the procedure for Generating self-signed root CA certificates using openssl in the section “Generating and Configuring a Root CA,” the correct command is **openssl req -new -x509 -days 1095 -key keys/ssl-proxy-ca.key -out certs/ssl-inspect-ca.cer**. Additionally, the **request security pki ca-certificate load ca-profile profile-ca1 filename profile-ca1.crt** has been added to Figure 1.

BGP Feature Guide for Security Devices

- In “Example: Configuring Route Authentication for BGP,” the following configuration steps in the CLI quick configuration and in the step-by-step procedure sections are not supported on SRX Series devices:

```
set security authentication-key-chains key-chain bgp-auth tolerance 30
set security authentication-key-chains key-chain bgp-auth key 0 secret
this-is-the-secret-password
set security authentication-key-chains key-chain bgp-auth key 0 start-time
2011-6-23.20:19:33-0700
set security authentication-key-chains key-chain bgp-auth key 1 secret
this-is-another-secret-password
set security authentication-key-chains key-chain bgp-auth key 1 start-time
2012-6-23.20:19:33-0700
```

Chassis Cluster Feature Guide for Security Devices

- In Step 5 of “Upgrading the Second Routing Engine When Using Chassis Cluster Dual Control Links on SRX5600 and SRX5800 Devices,” the bytes per second value is incorrectly shown as bs = 64k. The actual value is 1m.

- The **set chassis cluster cluster-id cluster-id node node reboot** operational mode command is missing from the Administration tab. This operational mode command sets the chassis cluster identifier (ID) and node ID on each device, and reboots the devices to enable clustering. This command has two options: **cluster-id cluster-id** (0 through 255) and **node node** (0 or 1). The system uses the chassis cluster ID and chassis cluster node ID to apply the correct configuration for each node (for example, when you use the **apply-groups** command to configure the chassis cluster management interface). The chassis cluster ID and node ID statements are written to the EPROM, and the statements take effect when the system is rebooted. Setting a cluster ID to 0 is equivalent to disabling a cluster. Support for extended cluster identifiers (more than 15 identifiers) added in Junos OS Release 12.1X46-D10. A cluster ID greater than 15 can only be set when the fabric and control link interfaces are connected back-to-back. The command has the following privilege level: maintenance.

If you have a cluster set up and running with an earlier release of Junos OS, you can upgrade to Junos OS Release 12.1X46-D10 or later and re-create a cluster with cluster IDs greater than 16. If for any reason you decide to revert to the previous version of Junos OS that did not support extended cluster IDs, the system comes up with standalone devices after you reboot. If the cluster ID set is less than 16 and you roll back to a previous release, the system comes back with the previous setup.

J-Web

- **J-Web pages for stateless firewall filters**—There is no documentation describing the J-Web pages for stateless firewall filters. To find these pages in J-Web, go to **Configure>Security>Firewall Filters**, and then select **IPv4 Firewall Filters** or **IPv6 Firewall Filters**. After configuring the filters, select **Assign to Interfaces** to assign your configured filters to interfaces.

Junos OS CLI User Guide

- In the **log-prefix** topic, SRX Series is missing from the list of supported platforms and release information.

SNMP MIBs and Traps Reference

- The “Enterprise-Specific MIBs and Supported Devices” topic incorrectly states that the SNMP IDP MIB is supported on high-end SRX Series devices. The SNMP IDP MIB is not supported on high-end SRX Series devices.

Modem Interfaces Feature Guide for Security Devices

- The Example: Configuring the 3G Wireless Modem Interface in Modem Interfaces Guide provides the following incorrect information for configuring a dialer filter for the 3G wireless modem interface:
 - `user@host# set firewall family inet dialer-filter corporate-traffic-only term term1 from source-address 20.20.90.4/32`
 - `user@host# set firewall family inet dialer-filter corporate-traffic-only term term1 from destination-address 200.200.201.1/32`

- user@host# **set firewall family inet dialer-filter corporate-traffic-only term term1 then note**

The following incorrect configuration output is included:

```
[edit]
user@host# show firewall family inet dialer-filter corporate-traffic-only
term term1 {
  from {
    source-address {
      20.20.90.4/32;
    }
    destination-address {
      200.200.201.1/32;
    }
  }
  then note;
}
```

The correct configuration is:

```
user@host# set firewall family inet dialer-filter corporate-traffic-only term term1 then
note
```

The following configuration is output from the correct configuration:

```
[edit]
user@host# show firewall
family inet {
  dialer-filter corporate-traffic-only {
    term term-1 {
      then note;
    }
  }
}
```

Multicast Feature Guide for Security Devices

- Multicast Source Discovery Protocol (MSDP) is not supported on SRX Series devices in any type of custom routing instance.

Network Address Translation

The command **show security nat source persistent-nat-table** under **Network Address Translation > Administration > Source NAT Operational Commands** has the following errors:

- The command is missing the **summary** option: **summary**—Display persistent NAT bindings summary.
- The command contains incomplete sample output —The corrected sample output is as follows:

show security nat source persistent-nat-table internal-ip internal-port

```
user@host> show security nat source persistent-nat-table internal-ip 9.9.9.1 internal-port 60784
```

```

Internal          Reflective      Source      Type
Left_time/ Curr_Sess_Num/ Source
In_IP In_Port I_Proto Ref_IP Ref_Port R_Proto NAT Pool
Conf_time Max_Sess_Num NAT Rule
9.9.9.1 60784 udp 66.66.66.68 60784 udp dynamic-customer-source
any-remote-host 254/300 0/30 105

```

show security nat source persistent-nat-table all

```

user@host> show security nat source persistent-nat-table all
Internal          Reflective      Source      Type
Left_time/ Curr_Sess_Num/ Source
In_IP In_Port I_Proto Ref_IP Ref_Port R_Proto NAT Pool
Conf_time Max_Sess_Num NAT Rule
9.9.9.1 63893 tcp 66.66.66.68 63893 tcp dynamic-customer-source
any-remote-host 192/300 0/30 105
9.9.9.1 64014 udp 66.66.66.68 64014 udp dynamic-customer-source
any-remote-host 244/300 0/30 105
9.9.9.1 60784 udp 66.66.66.68 60784 udp dynamic-customer-source
any-remote-host 254/300 0/30 105
9.9.9.1 57022 udp 66.66.66.68 57022 udp dynamic-customer-source
any-remote-host 264/300 0/30 105
9.9.9.1 53009 udp 66.66.66.68 53009 udp dynamic-customer-source
any-remote-host 268/300 0/30 105
9.9.9.1 49225 udp 66.66.66.68 49225 udp dynamic-customer-source
any-remote-host 272/300 0/30 105
9.9.9.1 52150 udp 66.66.66.68 52150 udp dynamic-customer-source
any-remote-host 274/300 0/30 105
9.9.9.1 59770 udp 66.66.66.68 59770 udp dynamic-customer-source
any-remote-host 278/300 0/30 105
9.9.9.1 61497 udp 66.66.66.68 61497 udp dynamic-customer-source
any-remote-host 282/300 0/30 105
9.9.9.1 56843 udp 66.66.66.68 56843 udp dynamic-customer-source
any-remote-host -/300 1/30 105

```

show security nat source persistent-nat-table summary

```

user@host> show security nat source persistent-nat-table summary
Persistent NAT Table Statistics on FPC5 PIC0:
binding total : 65536
binding in use : 0
enode total : 524288
enode in use : 0

```

Routing Protocols Overview for Security Devices

- The default route preference value in the “Understanding Route Preference Values” topic for Static and Static LSPs lists the values incorrectly. The correct values are as follows:

How Route Is Learned	Default Preference
Static	5
Static LSPs	6

Security Policy Applications Feature Guide for Security Devices

- The **show security policies** command output description is missing the definition for the following **Policy statistics** fields:
 - Output packets**—The total number of packets actually processed by the device.
 - Session rate**—The total number of active and deleted sessions.
- On the Overview tab, under IP-Related Predefined Policy Applications, in the topic entitled “Understanding IP-Related Predefined Policy Applications,” the Port column for both TCP-ANY and UDP-ANY should indicate 0-65535. The lead-in sentence should read, “Each entry includes the port and a description of the application.” TCP-ANY means any application that is using TCP, so there is no default port for it. The same is true for UDP-ANY.
- In the topic entitled “Understanding Miscellaneous Predefined Policy Applications,” table “Predefined Miscellaneous Applications” is incomplete. Under the RADIUS row, add a new row:

Table 18: Predefined Miscellaneous Applications

Application	Port	Description
RADIUS Accounting	1813	Enables the collecting of statistical data about users logging in to or out from a LAN and sending the data to a RADIUS Accounting server.

In table “Predefined Miscellaneous Applications” replace the IPsec-NAT row with the following:

Table 19: Predefined Miscellaneous Applications

Application	Port	Description
IKE	500	Internet Key Exchange is the protocol that sets up a security association in the IPsec protocol suite.
IKE-NAT	4500	Helps to perform Layer 3 NAT for S2C IKE traffic.

Table 20: Predefined Miscellaneous Applications

Application	Port	Description
VoIP	389	Internet Locator Service (ILS)
	522	User Location Service (ULS)
	1503	T.120 Data sharing
	1719	H.225 RAS message
	1720	Q.931 Call Setup
	1731	Audio Call Control
	5060	SIP protocol

Various Guides

- Some Junos OS user, reference, and configuration guides—for example the [Junos Software Routing Protocols Configuration Guide](#), [Junos OS CLI User Guide](#), and [Junos OS System Basics Configuration Guide](#)—mistakenly do not indicate SRX Series device support in the “Supported Platforms” list and other related support information; however, many of those documented Junos OS features are supported on SRX Series devices. For full, confirmed support information about SRX Series devices, please refer to Feature Explorer:
<http://pathfinder.juniper.net/feature-explorer/select-software.html?swName=Junos+OS&typ=1>.

Documentation Updates for the Junos OS Hardware Documentation

This section lists outstanding issues with the hardware documentation.

SRX5600 Services Gateway Hardware Guide

- The “Accessory Box Parts List” table in the “Verifying the SRX5600 Services Gateway Parts Received” topic lists the quantities for split washers, DC power terminal lugs, and 3 in. x 5 in. pink bag incorrectly. The correct quantities are as follows:

Part	Quantity
Split washers 1/4	34
DC power terminal lugs, 6-AWG	9
3 in. x 5 in. pink bag	5

The “Accessory Box Parts List” table in the “Verifying the SRX5600 Services Gateway Parts Received” topic is missing the following information:

Part	Quantity
Screws (4 x 8 mm long, 1.5 mm pitch)	4
SFP, Gigabit Ethernet, 850 nm, 550 m reach, SX, DDM	2
Fiber optic cable, Duplex, LC/LC, Multimode, 3 m	1

Related Documentation

- [New and Changed Features on page 95](#)
- [Known Behavior on page 135](#)
- [Known Issues on page 155](#)
- [Resolved Issues on page 156](#)

Migration, Upgrade, and Downgrade Instructions

This section includes the following topics:

- [Upgrading and Downgrading among Junos OS Releases on page 191](#)
- [Upgrading an AppSecure Device on page 193](#)
- [Network and Security Manager Support on page 194](#)
- [Upgrade and Downgrade Scripts for Address Book Configuration on page 194](#)
- [Upgrade Policy for Junos OS Extended End-Of-Life Releases on page 196](#)
- [Hardware Requirements on page 196](#)

Upgrading and Downgrading among Junos OS Releases

All Junos OS releases are listed in sequence on the JUNOS Software Dates & Milestones webpage:

<http://www.juniper.net/support/eol/junos.html>

To help in understanding the examples that are presented in this section, a portion of that table is replicated here. See [Table 8 on page 87](#). Note that releases footnoted with a 1 are Extended End-of-Life (EEOl) releases.

Table 21: Junos Software Dates & Milestones

Product	FRS Date
Junos 12.1X44 ¹²	01/18/2013
Junos 12.1	03/28/2012
Junos 11.4 ¹	12/21/2011
Junos 11.3	08/15/2011
Junos 11.2	08/03/2011
Junos 11.1	03/29/2011
Junos 10.4 ¹	12/08/2010
Junos 10.3	08/15/2010
Junos 10.2	05/28/2010
Junos 10.1	02/15/2010
Junos 10.0 ¹	11/04/2009
Junos 9.6	08/06/2009
Junos 9.5	04/14/2009
Junos 9.4	02/11/2009
Junos 9.3 ¹	11/14/2008
Junos 9.2	08/12/2008
Junos 9.1	04/28/2008
Junos 9.0	02/15/2008
Junos 8.5 ¹	fwd-srns-context

You can directly upgrade or downgrade between any two Junos OS releases that are within three releases of each other.

- Example: Direct release upgrade

Release 10.3 → (*bypassing Releases 10.4 and 11.1*) Release 11.2

To upgrade or downgrade between Junos OS releases that are more than three releases apart, you can upgrade or downgrade first to an intermediate release that is within three releases of the desired release, and then upgrade or downgrade from that release to the desired release.

- Example: Multistep release downgrade

Release 11.3 → (*bypassing Releases 11.2 and 11.1*) Release 10.4 → Release 10.3

Juniper Networks has also provided an even more efficient method of upgrading and downgrading using the Junos OS EEOL releases. EEOL releases generally occur once a calendar year and can be more than three releases apart. For a list of, EEOL releases, go to <http://www.juniper.net/support/eol/junos.html>

You can directly upgrade or downgrade between any two Junos OS EEOL releases that are within three EEOL releases of each other.

- Example: Direct EEOL release upgrade

Release 9.3 (EEOL) → (*bypassing Releases 10.0 [EEOL] and 10.4 [EEOL]*) Release 11.4 (EEOL)

To upgrade or downgrade between Junos OS EEOL releases that are more than three EEOL releases apart, you can upgrade first to an intermediate EEOL release that is within three EEOL releases of the desired EEOL release, and then upgrade from that EEOL release to the desired EEOL release.

- Example: Multistep release upgrade using intermediate EEOL release

Release 8.5 (EEOL) → (*bypassing Releases 9.3 [EEOL] and 10.0 [EEOL]*) Release 10.4 (EEOL) → Release 11.4 (EEOL)

You can even use a Junos OS EEOL release as an intermediate upgrade or downgrade step if your desired release is several releases later than your current release.

- Example: Multistep release upgrade using intermediate EEOL release

Release 9.6 → Release 10.0 (EEOL) → Release 10.2

For additional information about how to upgrade and downgrade, see the *Junos OS Installation and Upgrade Guide*.

Upgrading an AppSecure Device

Use the no-validate Option for AppSecure Devices.

For devices implementing AppSecure services, use the no-validate option when upgrading from Junos OS Release 11.2 or earlier to Junos OS 11.4R1 or later. The application signature package used with AppSecure services in previous releases has been moved from the configuration file to a signature database. This change in location can trigger an error

during the validation step and interrupt the Junos OS upgrade. The no-validate option bypasses this step.

[Network and Security Manager Support](#)

Network and Security Manager (NSM) support for High-End SRX Series Services Gateways with Junos OS 12.1X46-D10 is available only with NSM versions 2012.2R6 / 2012.1R10 and later. For additional information, see [Network and Security Manager](#) documentation.

[Upgrade and Downgrade Scripts for Address Book Configuration](#)

Beginning with Junos OS Release 11.4 or later, you can configure address books under the **[security]** hierarchy and attach security zones to them (zone-attached configuration). In Junos OS Release 11.1 and earlier, address books were defined under the **[security zones]** hierarchy (zone-defined configuration).

You can either define all address books under the **[security]** hierarchy in a zone-attached configuration format or under the **[security zones]** hierarchy in a zone-defined configuration format; the CLI displays an error and fails to commit the configuration if you configure both configuration formats on one system.

Juniper Networks provides Junos operation scripts that allow you to work in either of the address book configuration formats (see [Figure 7 on page 195](#)).

- [About Upgrade and Downgrade Scripts on page 194](#)
- [Running Upgrade and Downgrade Scripts on page 195](#)

About Upgrade and Downgrade Scripts

After downloading Junos OS Release 12.1, you have the following options for configuring the address book feature:

- **Use the default address book configuration**—You can configure address books using the zone-defined configuration format, which is available by default. For information on how to configure zone-defined address books, see the Junos OS Release 11.1 documentation.
- **Use the upgrade script**—You can run the upgrade script available on the Juniper Networks support site to configure address books using the new zone-attached configuration format. When upgrading, the system uses the zone names to create address books. For example, addresses in the trust zone are created in an address book named **trust-address-book** and are attached to the trust zone. IP prefixes used in NAT rules remain unaffected.

After upgrading to the zone-attached address book configuration:

- You cannot configure address books using the zone-defined address book configuration format; the CLI displays an error and fails to commit.
- You cannot configure address books using the J-Web interface.

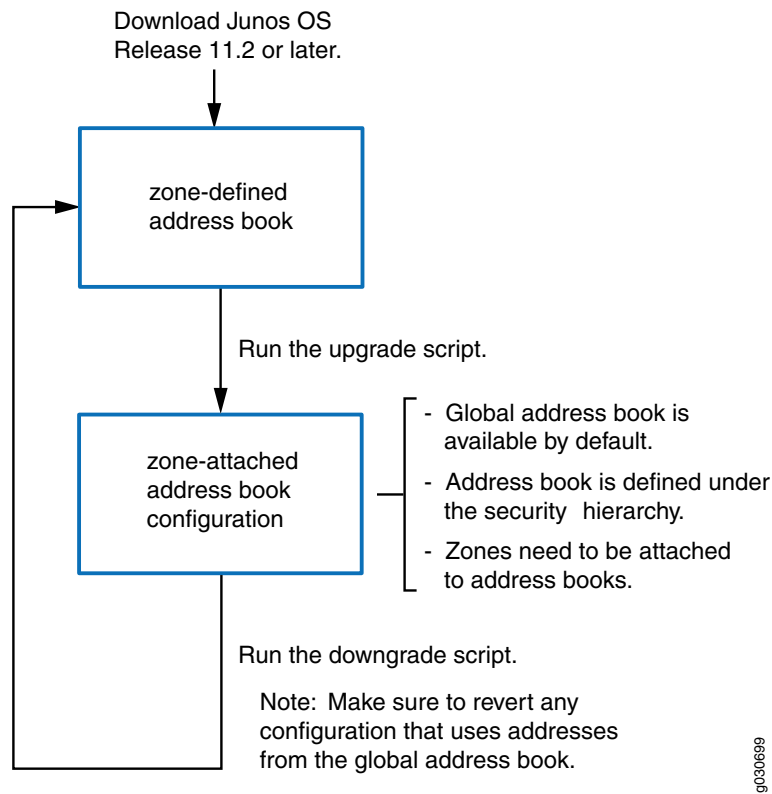
For information on how to configure zone-attached address books, see the Junos OS Release 11.4 documentation.

- **Use the downgrade script**—After upgrading to the zone-attached configuration, if you want to revert to the zone-defined configuration, use the downgrade script available on the Juniper Networks support site. For information on how to configure zone-defined address books, see the Junos OS Release 11.1 documentation.



NOTE: Before running the downgrade script, make sure to revert any configuration that uses addresses from the global address book.

Figure 7: Upgrade and Downgrade Scripts for Address Books



Running Upgrade and Downgrade Scripts

The following restrictions apply to the address book upgrade and downgrade scripts:

- The scripts cannot run unless the configuration on your system has been committed. Thus, if the zone-defined address book and zone-attached address book configurations are present on your system at the same time, the scripts will not run.
- The scripts cannot run when the global address book exists on your system.
- If you upgrade your device to Junos OS Release 11.4 or later and configure logical systems, the master logical system retains any previously configured zone-defined address book configuration. The master administrator can run the address book upgrade script to convert the existing zone-defined configuration to the zone-attached

configuration. The upgrade script converts all zone-defined configurations in the master logical system and user logical systems.



NOTE: You cannot run the downgrade script on logical systems.

For information about implementing and executing Junos operation scripts, see the *Junos OS Configuration and Operations Automation Guide*.

Upgrade Policy for Junos OS Extended End-Of-Life Releases

Support for upgrades and downgrades that span more than three Junos OS releases at a time is not provided, except for releases that are designated as Extended End-of-Life (EEOL) releases. EEOL releases provide direct upgrade and downgrade paths—you can upgrade directly from one EEOL release to the next EEOL release even though EEOL releases generally occur in increments beyond three releases.

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 10.0, 10.4, and 11.4 are EEOL releases. You can upgrade from Junos OS Release 10.0 to Release 10.4 or even from Junos OS Release 10.0 to Release 11.4. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind. For example, you cannot directly upgrade from Junos OS Release 10.3 (a non-EEOL release) to Junos OS Release 11.4 or directly downgrade from Junos OS Release 11.4 to Junos OS Release 10.3.

To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

For more information on EEOL releases and to review a list of EEOL releases, see <http://www.juniper.net/support/eol/junos.html>.

Hardware Requirements

Transceiver Compatibility for SRX Series Devices

We strongly recommend that only transceivers provided by Juniper Networks be used on high-end SRX Series Services Gateways interface modules. Different transceiver types (long-range, short-range, copper, and others) can be used together on multiport SFP interface modules as long as they are provided by Juniper Networks. We cannot guarantee that the interface module will operate correctly if third-party transceivers are used.

Please contact Juniper Networks for the correct transceiver part number for your device.

Related Documentation

- [New and Changed Features on page 95](#)
- [Documentation Updates on page 183](#)
- [Changes in Behavior and Syntax on page 113](#)

Product Compatibility

- [Hardware Compatibility on page 197](#)

Hardware Compatibility

To obtain information about the components that are supported on the device, and special compatibility guidelines with the release, see the SRX Series Hardware Guide.

To determine the features supported on SRX Series devices in Junos OS Release 12.1X46-D10, use the Juniper Networks Feature Explorer, a Web-based application that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. Find Feature Explorer at:

<http://pathfinder.juniper.net/feature-explorer/>.

Finding More Information

For the latest, most complete information about known and resolved issues with the Junos OS, see the Juniper Networks Problem Report Search application at:

<http://prsearch.juniper.net>.

Juniper Networks Feature Explorer is a Web-based application that helps you to explore and compare Junos OS feature information to find the correct software release and hardware platform for your network. Find Feature Explorer at:

<http://pathfinder.juniper.net/feature-explorer/>.

Juniper Networks Content Explorer is a Web-based application that helps you explore Juniper Networks technical documentation by product, task, and software release, and download documentation in PDF format. Find Content Explorer at:

<http://www.juniper.net/techpubs/content-applications/content-explorer/>.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can send your comments to techpubs-comments@juniper.net, or fill out the documentation feedback form at <https://www.juniper.net/cgi-bin/docbugreport/>. If you are using e-mail, be sure to include the following information with your comments:

- Document name
- Document part number
- Page number
- Software release version

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 postsales 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/customers/support/downloads/710059.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: <https://www.juniper.net/alerts/>
- 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 located at <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, visit us at <http://www.juniper.net/support/requesting-support.html>.

If you are reporting a hardware or software problem, issue the following command from the CLI before contacting support:

```
user@host> request support information | save filename
```

To provide a core file to Juniper Networks for analysis, compress the file with the **gzip** utility, rename the file to include your company name, and copy it to **ftp.juniper.net/pub/incoming**. Then send the filename, along with software version information (the output of the **show version** command) and the configuration, to **support@juniper.net**. For documentation issues, fill out the bug report form located at <https://www.juniper.net/cgi-bin/docbugreport/>.

Revision History

14 July 2015—Revision 2, Junos OS 12.1X46-D35 – High-End SRX Series, Branch SRX Series, and J Series.

Copyright © 2015, Juniper Networks, Inc. All rights reserved.

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

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