JUNIPER NETWORKS | Engineering Simplicity

IN FOCUS

Junos® OS Release 20.2

Published
2020-06-19

junos

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

IN FOCUS Junos® OS Release 20.2

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

1 **Start Here with Junos OS Release 20.2**

What You Need to Know About the In Focus Guide | 5

Important Features in Junos OS Release 20.2 | 5

2 **Protect and Reroute Layer 2 Services Using BGP Labeled Unicast**

How to Protect and Reroute Traffic Using BGP Labeled Unicast | 12

BGP PIC Edge Using BGP Labeled Unicast Overview | **12**

Benefits of BGP PIC Edge Using BGP Labeled Unicast | **12**

How does BGP Prefix Independent Convergence Work? | **13**

BGP PIC Edge Using BGP Labeled Unicast as the Transport Protocol Overview | **13**

3 **Implement Retaining the Authentication Session Using IP-MAC Bindings**

How to Implement Retaining the Authentication Session Using IP-MAC Bindings | 16

Retaining the Authentication Session Based on IP-MAC Address Bindings | **16**

Benefits | **17**

CLI Configuration | **17**

RADIUS Server Attributes | **18**

Verification | **18**

4 **Analyze Unknown Application Traffic Using Packet Capture**

How to Configure Packet Capture for Unknown Application Traffic | 21

Packet Capture of Unknown Application Traffic Overview | **21**

Benefits of Packet Capture of Unknown Application Traffic | **21**

Configure Packet Capture of Unknown Application Traffic | **22**

5 **Control the Re-merge Behavior on Point-to-Multipoint LSP Network**

How to Control the Re-merge Behavior on the Point-to-Multipoint LSP Network | 30

Re-merge Behavior on Point-to-Multipoint LSP Overview | **30**

Benefits of Controlling the P2MP LSP Re-merge | **30**

What is P2MP LSP Re-merge? | **31**

Modify the Default P2MP LSP Re-merge Behavior | **32**

CHAPTER 1

Start Here with Junos OS Release 20.2

What You Need to Know About the In Focus Guide | 5

Important Features in Junos OS Release 20.2 | 5

What You Need to Know About the In Focus Guide

Use this guide to quickly learn about the most important features in Junos OS Release 20.2 and how you can deploy them in your network.

You might also be interested in seeing the complete list of features in the Release Notes for Junos OS Release 20.2. In addition to this guide, you can find detailed information on concepts, configuration, and examples in the Junos OS documentation.

Want to tell us what you think about this guide? E-mail us at techpubs-comments@juniper.net.

Important Features in Junos OS Release 20.2

For details on these features, go to the other chapters in this guide or click the link in the feature description below.

- **Cloud-init support for Azure (vSRX 3.0)**—Starting in Junos OS Release 20.2R1, vSRX 3.0 on Azure supports cloud-init for VM provisioning. vSRX 3.0 calls the imported cloud-init package (version 18.4) during the initial boot process and applies specified user-data files to the new configurations. You can choose cloud-init to apply the required configurations during deployment. When you use cloud-init for configurations, after the VM boots up, you can easily configure the VM in few steps.

  [See Using Cloud-Init to Automate the Initialization of vSRX 3.0 Instances in Microsoft Azure Cloud.]

- **MX Series Virtual Chassis support for the ephemeral database (MX Series)**—Starting in Junos OS Release 20.2R1, MX Series Virtual Chassis support configuring the ephemeral database. The ephemeral database is an alternate configuration database that provides a fast programmatic interface for performing configuration updates on devices running Junos OS.

  [See Understanding the Ephemeral Configuration Database.]

- **Retain the authentication session based on DHCP or SLAAC snooping entries (EX2300, EX3400, and EX4300)**—Starting in Junos OS Release 20.2R1, you can configure the switching device to check for a DHCP, DHCPv6, or SLAAC snooping entry before terminating the authentication session when the MAC address ages out. If a snooping entry is present, the authentication session for the end device with that MAC address remains active. This ensures that the end device will be reachable even if the MAC address ages out.

  [See “How to Implement Retaining the Authentication Session Using IP-MAC Bindings” on page 16.]

- **Rest API support for EX Series switches**—Starting in Release 20.2R1, Junos OS supports all REST API features except HTTP sessions on EX Series switches. The REST API enables you to securely connect to the Junos OS devices, execute remote procedure calls (RPC) commands, use REST API explorer GUI

to conveniently experiment with any of the REST APIs, and use a variety of formatting and display options including JavaScript Object Notation (JSON).

[See REST API Guide.]

- **CPU usage monitoring (SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 20.2R1, you can use the following operational commands to monitor the average CPU usage information for the last minute, hour, or day of an SPC3 card:
  - **show security monitoring performance spu summary fpc *fpc-slot-number* pic *pic-slot-number***
  - **show security monitoring performance spu summary fpc *fpc-slot-number* pic *pic-slot-number* thread *thread-number***

  You can monitor the CPU usage information only when the PIC is online.

  We've introduced the new SNMP MIBs **jnxJsSPUMonitoringSPUThreadsNumber**, **jnxJsSPUMonitoringSPUThreadIndex**, **jnxJsSPUMonitoringSPUThreadLastMinUsage**, **jnxJsSPUMonitoringSPUThreadLastHourUsage**, and **jnxJsSPUMonitoringSPUThreadLastDayUsage** to monitor the CPU usage information of an SPC3 card.

  [See show snmp mib and show security monitoring performance spu.]

- **Contrail networking support (cSRX)**—Starting in Junos OS Release 20.2R1, we have integrated cSRX Container Firewall into a Contrail network as a distributed host-based firewall service on a Docker container. Using this deployment, you can obtain agile, elastic, and cost-saving security services.

  The new virtual solution provides the following capabilities:

  - Layer 7 security protection (antivirus, application firewall, IPS, application identification, URL filtering, user firewall, UTM content and Web filtering only)
  - Automated service provisioning and orchestration
  - Distributed and multitenant traffic securing
  - Centralized management with Junos Space Security Director, including dynamic policy and address update, remote log collections, and security events monitoring
  - Scalable security services with small footprints

  [See cSRX as Contrail Host-based Firewall User Guide.]

- **Safe search enhancement for Web filtering (SRX Series and vSRX)**—Starting in Junos OS Release 20.2R1, we've introduced safe search UTM Web filtering on well-known search engines. This safe search enhancement enforces the safest Web browsing mode available, by default. You can disable the safe search option at the Web filtering-level and profile-level configurations. You can also block search engine cache on the well-known search engines. By blocking the search engine cache, you can hide your Web-browsing activities from other users if you are a part of an organization that has multiple Web users in educational, financial, health-care, banking, and corporate segments.

  [See Safe Search Enhancement for Web Filtering.]

- **Packet capture of unknown application traffic (NFX Series, SRX Series, and vSRX)**—Starting in Junos OS Release 20.2R1, we've added new capability to your security device that allows you to capture unknown application traffic.

  Once you have configured the packet capture options on your security device, the unknown application traffic information is gathered and stored on the device in a packet capture file (**.pcap**). You can use the packet capture of an unknown application to define a new custom application signature. You can use this custom application signature in a security policy to manage the application traffic more efficiently.

  You can also send the **.pcap** file to Juniper Networks in case where the traffic is incorrectly classified, or to request for the creation of an application signature.

  [See "How to Configure Packet Capture for Unknown Application Traffic" on page 21 and Application Identification.]

- **Support for Must-IE check and IE removal for GTPv1 and GTPv2 (SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)**—Starting in Release 20.2R1, Junos OS supports the following information element (IE) enforcement functions for GTPv1 and GTPv2:

  - **Must-IE check**: Use this function to check the presence of IEs in GTPv1-C and GTPv2-C messages, which helps to verify message integrity. The device checks the presence of Must-IEs of specific GTP messages and forwards the messages only if Must-IEs are present.
  - **IE removal**: Use this function to remove IEs from GTPv1-C and GTPv2-C. This helps to retain interoperability between Second-Generation Partnership Project (2GPP) and Third-Generation Partnership Project (3GPP) networks.

  [See Example: Configure Must-IE check for GTPv1 and GTPv2, and Example: Configure IE removal for GTPV1 and GTPv2.]

- **User authentication support for tenant systems (SRX Series)**—Starting in Release 20.2R1, Junos OS introduces the following authentication support for tenant systems:

  - **address-assignment pools**: Creates centralized IPv4 and IPv6 address pools independent of the client applications that use the pools.
  - **access profiles**: Runs authentication and accounting requests.
  - **clear network-access aaa subscribers**: Clears AAA subscriber statistics and logs out subscribers. You can log out subscribers based on the username or on the subscriber session identifier.

  [See Firewall Authentication for Tenant Systems.]

- **TI-LFA SRLG protection for IS-IS (MX Series and PTX Series)**—Starting in Junos OS Release 20.2R1, you can configure Shared Risk Link Group (SRLG) protection in topology-independent loop-free alternate (TI-LFA) networks for segment routing. IS-IS computes the fast reroute path that is aligned with the post-convergence path and excludes the SRLG of the protected link. All local and remote links that share any SRLG with the protecting link are excluded. The point of local repair (PLR) sets up the label stack for the fast reroute path with a different outgoing interface.

To enable TI-LFA SRLG protection with segment routing for IS-IS, include the **srlg-protection** statement at the **[edit protocols isis interface *name* level *number* post-convergence-lfa]** hierarchy level.

[See Understanding Topology-Independent Loop-Free Alternate with Segment Routing for IS-IS.]

- **Support for Layer 2 circuit, Layer 2 VPN, and VPLS services with BGP labeled unicast (MX Series, EX9204, EX9208, EX9214, EX9251, and EX9253 devices)**—Starting with Junos OS Release 20.2R1, MX Series, EX9204, EX9208, EX9214, EX9251, and EX9253 devices support Layer 2 circuit, Layer 2 VPN, and VPLS services with BGP labeled unicast as the transport protocol. BGP PIC Edge using the BGP labeled unicast transport protocol helps to protect traffic failures over border nodes (ABR and ASBR) in multi-domain networks. Multi-domain networks are typically used in metro-aggregation and mobile backhaul networks designs.

  To enable expanded hierarchical nexthop structure for Layer 2 services (Layer 2 circuit, Layer 2 VPN, and VPLS services), you need to configure the following CLI configuration statements at the [**edit protocols**] hierarchy level:

  Layer 2 circuit:

  ```
  [edit protocols]
  user@host#set l2circuit resolution preserve-expanded-heirarchy;
  ```

  Layer 2 VPN:

  ```
  [edit protocols]
  user@host#set l2vpn resolution preserve-expanded-heirarchy;
  ```

  [See Load Balancing for a BGP Session.]

- **Control the default re-merge behavior on the P2MP LSP (MX Series)**—Starting with Junos OS Release 20.2R1, you can control and change the default re-merge behavior on RSVP P2MP LSP. The term re-merge refers to the case of an ingress (headend) or transit node (re-merge node) that creates a re-merge branch intersecting the P2MP LSP at another node in the network. This may occur due to events such as an error in path calculation, an error in manual configuration, or network topology changes during the establishment of the P2MP LSP.

  You can control the default re-merge behavior on P2MP LSPs by enabling the newly introduced **no-re-merge** and **no-p2mp-re-merge** CLI commands at the ingress (headend) and transit devices (re-merge nodes), respectively.

  [See Re-merge Behavior on Point-to-Multipoint LSP Overview.]

- **Support for security feeds in security policies (SRX Series and vSRX)**—Starting in Junos OS Release 20.2R1, you can add source and destination addresses to the security intelligence (SecIntel) profiles to generate security feeds in a security policy. You can accomplice this by configuring the **security-intelligence** command. After the feeds are generated, you can configure other security policies to use the feeds as a **dynamic-address** to match designated traffic and perform policy actions.

You can configure the **security-intelligence** command as permit, deny, or reject match conditions in a security policy at the following hierarchy levels:

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit
    application-services]
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then deny
    application-services]
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then reject
    application-services]
```

[See security-intelligence and Encrypted Traffic Analysis Overview*Encrypted Traffic Analysis Overview*.]

- **Support for BGP-LU over SR-TE and IS-IS segment routing underlay**—Starting in Junos OS Release 20.2R1, BGP Labeled Unicast (BGP-LU) can resolve IPv4 or IPv6 routes over segment routing–traffic engineering (SR-TE) with IS-IS underlay for both IPv4 and IPv6 address families. BGP-LU supports mapping a BGP community color and a resolution map for SR-TE. A colored protocol next hop is constructed and it is resolved on a colored SR-TE tunnel in the **inetcolor.0** or **inet6color.0** table. Thus BGP-LU resolves protocol next hop over SR-TE tunnels for packet transport.

  [See Color-Based Mapping of VPN Services Overview.]

- **Increased port block allocation size (SRX5000 line of devices with SPC2 and SPC3 cards)**—we've increased the port block allocation size so you can store more log files in the log server.

  - When you disable **interim log**, you can increase the size of port block allocation from *64* to *8* .
  - When you enable **interim log**, you can increase the size of port block allocation from *128* to *8*.

  If you configure the port block allocation size less than 8, the system displays the warning message **warning: To save system memory, the block size is recommended to be no less than 8**.

  [See Guidelines for Configuring Secured Port Block Allocation and Configure Port Block Allocation Size.]

- **VMware Tools support for VMware Hypervisors (vSRX 3.0)**—Starting in Junos OS Release 20.2R1, vSRX 3.0 on VMware Hypervisors support VMware Tools version 10.3.0 for autoconfiguration. The VMware Tools are initialized when the guest operating system starts. The service passes information between the host and guest operating systems for better management and operation.

  [See Automate the Initialization of vSRX 3.0 Instances on VMware Hypervisor using VMware Tools.]

- **Policy-based threat profile for IDP (SRX Series)**—Starting from Junos OS Release 20.2R1, you can configure IDP rules with threat profiles to define attacker IP and target IP feeds.

  When traffic matches the feed data, IDP provides feed update to add the IP information in the Security Intelligence (SecIntel) module.

  This feature allows the SRX Series device to identify threats, and propagate intelligence for real-time enforcement and provides the ability to perform endpoint classification.

[See IDP Policy Rules and IDP Rule Bases, security-intelligence, and Encrypted Traffic Analysis Overview.]

2

CHAPTER

Protect and Reroute Layer 2 Services Using BGP Labeled Unicast

How to Protect and Reroute Traffic Using BGP Labeled Unicast | 12

How to Protect and Reroute Traffic Using BGP Labeled Unicast

SUMMARY

Learn how to protect and reroute traffic to the destination network if a border node fails in a multi-domain network.

IN THIS SECTION

- BGP PIC Edge Using BGP Labeled Unicast Overview | **12**

BGP PIC Edge Using BGP Labeled Unicast Overview

SUMMARY

This section talks about the benefits and overview of BGP PIC Edge using BGP labeled unicast as the transport protocol.

IN THIS SECTION

- Benefits of BGP PIC Edge Using BGP Labeled Unicast | **12**
- How does BGP Prefix Independent Convergence Work? | **13**
- BGP PIC Edge Using BGP Labeled Unicast as the Transport Protocol Overview | **13**

Benefits of BGP PIC Edge Using BGP Labeled Unicast

- Provides traffic protection in case of border (ABR and ASBR) node failures in multi-domain networks.
- Provides faster restoration of network connectivity and reduces traffic loss if the primary path becomes unavailable.

How does BGP Prefix Independent Convergence Work?

BGP Prefix Independent Convergence (PIC) improves BGP convergence on network node failures. BGP PIC creates and stores primary and backup paths for the indirect next hop on the Routing Engine and also provides the indirect next hop route information to the Packet Forwarding Engine. When a network node failure occurs, the Routing Engine signals the Packet Forwarding Engine that an indirect next hop has failed, and that the traffic is rerouted to a pre-calculated equal-cost or backup path without modifying BGP prefixes. Routing the traffic to the destination prefix continues by using the backup path to reduce traffic loss until the global convergence through BGP is resolved.

BGP convergence is applicable to both core and edge network node failures. In the case of BGP PIC Core, adjustments to the forwarding chains are made as a result of node or core link failures. In the case of BGP PIC Edge, adjustments to the forwarding chains are made as a result of edge node or edge link failures.

BGP PIC Edge Using BGP Labeled Unicast as the Transport Protocol Overview

BGP PIC Edge using the BGP labeled unicast transport protocol helps to protect and reroute traffic when border nodes (ABR and ASBR) failures happen in multi-domain networks. Multi-domain networks are typically used in Metro Ethernet aggregation and mobile backhaul network designs.

On Juniper Networks MX Series, EX Series, and PTX Series routers, BGP PIC Edge supports IPv4, IPv6, BGP labeled unicast, Layer 3 VPN, Layer 2 VPN, Layer 2 circuit, and VPLS services. These BGP services are multipath (learnt from multiple PEs) and resolved through BGP labeled unicast routes, which could again be a multipath learnt from other ABRs. Transport protocols supported over BGP PIC Edge are RSVP, LDP, OSPF, and ISIS. Starting from Junos OS Release 20.2R1, MX Series, EX9204, EX9208, EX9214, EX9251, and EX9253 devices supports the Layer 2 services such as Layer 2 circuit, Layer 2 VPN, and VPLS with BGP labeled unicast.

The following BGP service routes are supported:

- IPv6 Layer 3 VPN services over IPv4 BGP labeled unicast
- IPv4 services over IPv4 BGP labeled unicast
- IPv4 Layer 3 VPN services over IPv4 BGP labeled unicast
- IPv6 BGP labeled unicast service over IPv4 BGP labeled unicast

EX Series, MX Series, and PTX Series routers support BGP PIC Edge with BGP labeled unicast as the transport protocol.

WHAT'S NEXT

> For an example on configuring BGP PIC Edge using BGP labeled unicast, see Example: Protecting IPv4 Traffic over Layer 3 VPN Running BGP Labeled Unicast.

> For more information on configuring BGP PIC and load balancing, see Load Balancing for a BGP Session.

Release History Table

Release	Description
20.2R1	Starting from Junos OS Release 20.2R1, MX Series, EX9204, EX9208, EX9214, EX9251, and EX9253 devices supports the Layer 2 services such as Layer 2 circuit, Layer 2 VPN, and VPLS with BGP labeled unicast.

3

CHAPTER

Implement Retaining the Authentication Session Using IP-MAC Bindings

How to Implement Retaining the Authentication Session Using IP-MAC Bindings | **16**

How to Implement Retaining the Authentication Session Using IP-MAC Bindings

SUMMARY

You can prevent the authentication session for an end device from being terminated when the MAC address for that device ages out.

IN THIS SECTION

- Retaining the Authentication Session Based on IP-MAC Address Bindings | **16**

Retaining the Authentication Session Based on IP-MAC Address Bindings

IN THIS SECTION

- Benefits | **17**
- CLI Configuration | **17**
- RADIUS Server Attributes | **18**
- Verification | **18**

MAC RADIUS authentication is often used to permit hosts that are not enabled for 802.1X authentication to access the LAN. End devices such as printers are not very active on the network. If the MAC address associated with an end device ages out due to inactivity, the MAC address is cleared from the Ethernet switching table, and the authentication session ends. This means that other devices will not be able to reach the end device when necessary.

If the MAC address that ages out is associated with an IP address in the DHCP, DHCPv6, or SLAAC snooping table, that MAC-IP address binding will be cleared from the table. This can result in dropped traffic when the DHCP client tries to renew its lease.

You can configure the switching device to check for an IP-MAC address binding in the DHCP, DHCPv6, or SLAAC snooping table before terminating the authentication session when the MAC address ages out.

If the MAC address for the end device is bound to an IP address, then it will be retained in the Ethernet switching table, and the authentication session will remain active.

This feature can be configured globally for all authenticated sessions using the CLI, or on a per-session basis using RADIUS attributes.

Benefits

This feature provides the following benefits:

- Ensures that an end device is reachable by other devices on the network even if the MAC address ages out.
- Prevents traffic from dropping when the end device tries to renew its DHCP lease.

CLI Configuration

Before you can configure this feature:

- DHCP snooping, DHCPv6 snooping, or SLAAC snooping must be enabled on the device.
- The **no-mac-table-binding** CLI statement must be configured. This disassociates the authentication session table from the Ethernet switching table, so that when a MAC address ages out, the authentication session will be extended until the next reauthentication.

```
[edit]
user@switch# set protocols dot1x authenticator no-mac-table-binding;
```

To configure this feature globally for all authenticated sessions:

- Configure the switching device to check for an IP-MAC address binding in the DHCP, DHCPv6, or SLAAC snooping table before terminating the authentication session when the MAC address ages out using the **ip-mac-session-binding** CLI statement:

```
[edit]
user@switch# set protocols dot1x authenticator ip-mac-session-binding;
```

> **NOTE:** You cannot commit the **ip-mac-session-binding** configuration unless the **no-mac-table-binding** is also configured.

RADIUS Server Attributes

You can configure this feature for a specific authentication session using RADIUS server attributes. RADIUS server attributes are clear-text fields encapsulated in Access-Accept messages sent from the authentication server to the switching device when a supplicant connected to the switch is successfully authenticated.

To retain the authentication session based on IP-MAC address bindings, configure both of the following attribute-value pairs on the RADIUS server:

- Juniper-AV-Pair = “IP-Mac-Session-Binding”
- Juniper-AV-Pair = “No-Mac-Binding-Reauth”

The Juniper-AV-Pair attribute is a Juniper Networks vendor-specific attribute (VSA). Verify that the Juniper dictionary is loaded on the RADIUS server and includes the Juniper-AV-Pair VSA (ID# 52).

If you need to add the attribute to the dictionary, locate the dictionary file (**juniper.dct**) on the RADIUS server and add the following text to the file:

```
ATTRIBUTE Juniper-AV-Pair  Juniper-VSA(52, string) r
```

> **NOTE:** For specific information about configuring your RADIUS server, consult the AAA documentation included with your server.

Verification

Verify the configuration by issuing the operational mode command **show dot1x interface *interface-name* detail** and confirm that the **Ip Mac Session Binding** and **No Mac Session Binding** output fields indicate that the feature is enabled.

`user@switch>` **show dot1x interface ge-0/0/16.0 detail**

```
ge-0/0/16.0
 Role: Authenticator
 Administrative state: Auto
 Supplicant mode: Multiple
 Number of retries: 3
 Quiet period: 60 seconds
 Transmit period: 5 seconds
```

```
Mac Radius: Enabled
Mac Radius Restrict: Disabled
Mac Radius Authentication Protocol: EAP-MD5
Reauthentication: Disabled
Configured Reauthentication interval: 3600 seconds
Supplicant timeout: 30 seconds
Server timeout: 30 seconds
Maximum EAPOL requests: 2
Guest VLAN member: <not configured>
No Mac Session Binding: Enabled
Ip Mac Session Binding: Enabled
Number of connected supplicants: 1
  Supplicant: abc, 00:30:48:8C:66:BD
    Operational state: Authenticated
    Backend Authentication state: Idle
    Authentication method: Mac Radius
    Authenticated VLAN: v100
    Session Reauth interval: 3600 seconds
    Reauthentication due in 0 seconds
    Ip Mac Session Binding: Enabled
    No Mac Binding Reauth: Enabled
    Eapol-Block: Not In Effect
```

Clients authenticated with MAC RADIUS should remain authenticated, and MAC address entries in the Ethernet switching table should also be retained after expiration of the MAC timer.

WHAT'S NEXT

> *Authentication Session Timeout*

4

CHAPTER

Analyze Unknown Application Traffic Using Packet Capture

How to Configure Packet Capture for Unknown Application Traffic | 21

How to Configure Packet Capture for Unknown Application Traffic

SUMMARY

Packet capture for unknown application traffic feature captures the packet details and stores the information in a packet capture file (**.pcap**) on your security device. You can use the packet capture information to analyze the application traffic and gain more insight on unknown applications. You can use this information to define a new custom application signature to manage the application traffic more effectively.

IN THIS SECTION

- Packet Capture of Unknown Application Traffic Overview | **21**
- Configure Packet Capture of Unknown Application Traffic | **22**

Packet Capture of Unknown Application Traffic Overview

Unknown application traffic is the traffic that does not match an application signature.

You can use the packet capture of unknown applications functionality to gather more details about an unknown application on your security device.

Once you've configured packet capture options on your security device, the unknown application traffic is gathered and stored on the device in a packet capture file (**.pcap**). You can use the packet capture of an unknown application to define a new custom application signature. You can use this custom application signature in a security policy to manage the application traffic more efficiently.

You can send the **.pcap** file to Juniper Networks in case where the traffic is incorrectly classified, or to request for the creation of an application signature.

Benefits of Packet Capture of Unknown Application Traffic

You can use the packet capture of unknown application traffic to:

- Gather more insight about an unknown application
- Analyze unknown application traffic for potential threats

- Assist in creation of security policy rules
- Enable custom application signature creation

> **NOTE:** Implementing security policies that block all unknown application traffic could cause issues with network based applications. Before applying these types of policies, be sure to validate that this approach does not cause issues in your environment. You must carefully analyze the unknown application traffic, and accordingly you can define the security polices.

Configure Packet Capture of Unknown Application Traffic

IN THIS SECTION

- Before You Begin | 22
- Overview | 22
- Configuration | 23
- Verification | 27

In this example, you'll enable packet capture of unknown application traffic.

Before You Begin

To enable packet capture of unknown application traffic, you must:

- Install a valid application identification feature license on your SRX Series device. See Managing Junos OS Licenses.
- Download and install the Junos OS application signature package. Download and Install Junos OS Application Signature Package.
- Ensure you have Junos OS Release 20.2R1 or later version on your security device.

Overview

In this example, you'll learn how to configure the packet capture for unknown applications on your security device by completing the following steps:

- Define Packet Capture Options on page 23
- Set Packet Capture Mode on page 24
- Enable Packet Capture Globally or at Policy Level on page 24
- Access Packet Capture File on page 25

Configuration

In this example, you enable packet capturing on your security device by defining packet capture options such as maximum packet limit, maximum byte limit, and number of files. To learn about the packet capture configuration options, see *packet-capture* before you start the configuration

Define Packet Capture Options

Step-by-Step Procedure

To set the packet capture options:

1. Set the maximum number of UDP packets per session.

```
[edit]
user@host# set services application-identification packet-capture max-packets 10
```

2. Set the maximum number of TCP bytes per session.

```
[edit]
user@host# set services application-identification packet-capture max-bytes 2048
```

3. Set the maximum number of packet capture files to be created before the oldest file is overwritten by a new file which is known as file rotate frequency.

```
[edit]
user@host# set services application-identification packet-capture max-files 30
```

Set Packet Capture Mode

In this step, you set the packet capture mode.

You can capture the packets for the unknown application traffic in either of the following modes:

- ASC mode—Capture the packet for unknown application when the application is classified as junos:UNKNOWN and has a matching entry in the application system cache (ASC). This mode is enabled by default.
- Aggressive mode—Capture all traffic before AppID classifies the applications. In this mode, the system captures all application traffic irrespective of an available ASC entry. Packet capture starts for the first packet of the first session. Note that, aggressive mode is a more resource-intensive mode.

  To enable aggressive mode, use the following command:

  ```
  [edit]
  user@host# set services application-identification packet-capture aggressive-mode
  ```

  We do not recommend using aggressive mode unless you need to capture the first occurrence of a flow prior. The default behavior of the device relies on the ASC.

Enable Packet Capture Globally or at Policy Level

You can configure the packet capture globally to capture all unknown application traffic or enable the packet capture for application traffic specific to a security policy. In this example, you'll enable packet capture for unknown application traffic at the security policy-level.

Step-by-Step Procedure

- Configure packet capture of unknown application traffic that matches the security policy P1 rules.

  ```
  [edit]
  user@host# set security policies from-zone untrust to-zone trust policy P1 match source-address any
  user@host# set security policies from-zone untrust to-zone trust policy P1 match destination-address any
  user@host# set security policies from-zone untrust to-zone trust policy P1 match application any
  user@host# set security policies from-zone untrust to-zone trust policy P1 match dynamic-application
     junos:UNKNOWN
  user@host# set security policies from-zone untrust to-zone trust policy P1 then permit application-services
     packet-capture
  ```

  To enable packet capture for unknown application traffic at the security policy level, you must include **junos:UNKNOWN** as the dynamic-application match conditions.

  When you configure the security policy (P1), the system captures the packet details for the application traffic that matches the security policy match criteria.

Results

From configuration mode, confirm your configuration by entering the **[edit services application-identification]** and **[edit security policies]** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit services application-identification]
user@host# show
packet-capture {
    max-packets 10;
    max-bytes 2048;
    max-files 30;
}
```

```
[edit security policies]
user@host# show
from-zone untrust to-zone trust {
    policy P1 {
        match {
            source-address any;
            destination-address any;
            application any;
            dynamic-application [ junos:UNKNOWN ];
        }
        then {
            permit {
                application-services {
                    packet-capture;
                }
            }
        }
    }
}
```

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

Access Packet Capture File

After you complete the configuration and commit it, you can view the packet capture (**.pcap**) file. The system generates a unique packet capture file for each destination IP address, destination port, and protocol.

Step-by-Step Procedure

To view the packet capture file:

1. Navigate to the directory where packet capture files are stored on the device.

```
user@host> start shell
```

```
%
% cd /var/tmp/pcap
```

2. Locate the packet capture file. For example, locate the **5.0.0.1_3548_6.pcap** file.

   The packet capture file is saved in *destination-IP-address. destination-port.protocol. pcap* format.

   You can download the packet capture file by using SFTP or SCP and read the packet capture file by using Wireshark or any other packet capture reader tools.

   Figure 1 on page 26 shows a sample packet capture file generated for the unknown application traffic.

Figure 1: Sample Packet Capture File

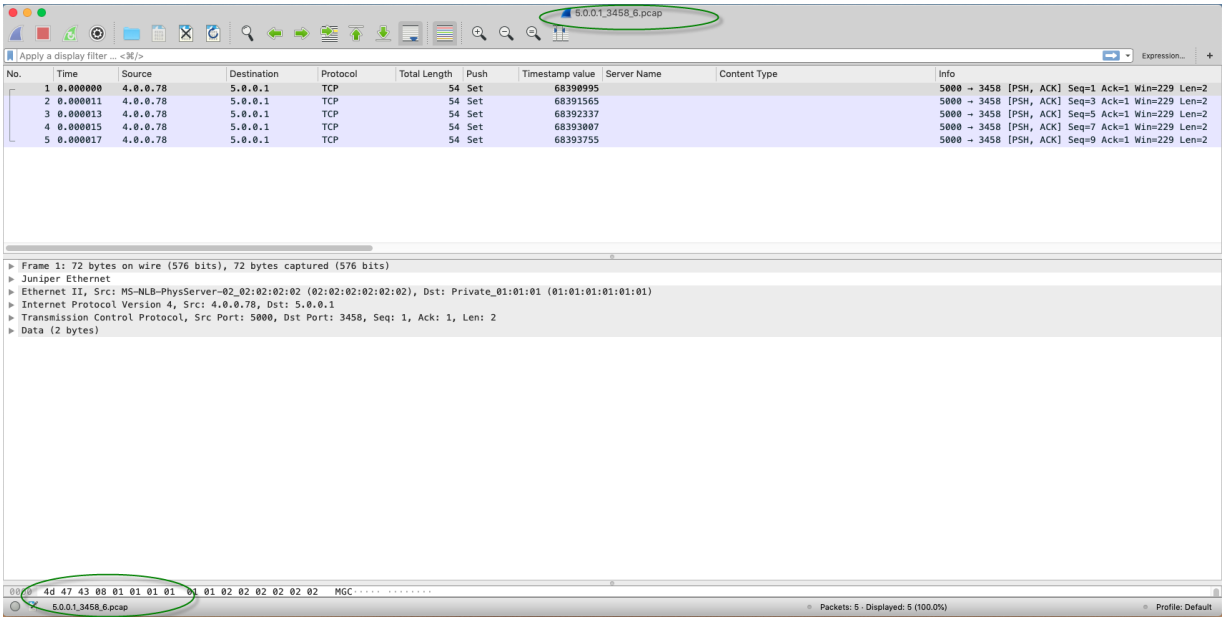

> **NOTE:** In some situations network traffic drops could cause the device to be unable to capture all packets. In this case, the **.pcap** file will reflect the missing packets.

The security device saves the packet capture details for all the traffic matching the three matching criteria (destination IP address, destination port, and protocol) in the same file irrespective of global or policy-level configuration. The system maintains the cache with the destination IP address, destination port, and the protocol and does not accept the repeated capturing of the same traffic more than the allowed limit. You can set the packet capture file options mentioned in *packet-capture*.

Verification

Viewing Packet Capture Details

Purpose

View the packet capture details to confirm that your configuration is working.

Action

Use the **show services application-identification packet-capture counters** command.

user@host> show services application-identification packet-capture counters

```
pic: 0/0
  Counter type                                                  Value
Total sessions captured                                             1
Total packets captured                                              6
Active sessions being captured                                      0
Sessions ignored because of memory allocation failures              0
Packets ignored because of memory allocation failures               0
Ipc messages ignored because of storage limit                       0
Sessions ignored because of buffer-packets limit                    0
Packets ignored because of buffer-packets limit                     0
Inconclusive sessions captured                                      0
Inconclusive sessions ignored                                       0
Cache entries timed out                                             0
```

Meaning

From this sample output, you can get details such as the number of sessions being captured, and the number of sessions already captured. For more details about the packet capture counters, see *show services application-identification packet-capture counters*.

SEE ALSO

packet-capture
show services application-identification packet-capture counters
request services application-identification clear packet-capture all
clear services application-identification packet-capture counters

WHAT'S NEXT

> For more information on application identification, see Application Identification. For details about custom applications, see Custom Application Signatures for Application Identification

5

CHAPTER

Control the Re-merge Behavior on Point-to-Multipoint LSP Network

How to Control the Re-merge Behavior on the Point-to-Multipoint LSP Network | 30

How to Control the Re-merge Behavior on the Point-to-Multipoint LSP Network

SUMMARY

Learn how to control and change the default P2MP sub LSP re-merge behavior in a P2MP RSVP MPLS network.

IN THIS SECTION

- Re-merge Behavior on Point-to-Multipoint LSP Overview | **30**

Re-merge Behavior on Point-to-Multipoint LSP Overview

SUMMARY

This section talks about the benefits and overview of controlling the re-merge behavior on RSVP Point-to-Multipoint (P2MP) LSPs.

IN THIS SECTION

- Benefits of Controlling the P2MP LSP Re-merge | **30**
- What is P2MP LSP Re-merge? | **31**
- Modify the Default P2MP LSP Re-merge Behavior | **32**

Benefits of Controlling the P2MP LSP Re-merge

- Reduces the RSVP signalling load on the ingress (headend routers) by avoiding path computation of sub LSPs which creates a re-merge condition.
- Saves the network bandwidth by rejecting the P2MP sub LSP re-merge at the transit node.

What is P2MP LSP Re-merge?

In a P2MP MPLS LSP network, the term re-merge refers to the case of an ingress (headend) or transit node (re-merge node) that creates a re-merge branch intersecting the P2MP LSP at another node down the tree. This may occur due to events such as an error in path calculation, an error in manual configuration, or network topology changes during the establishment of the P2MP LSP.

RFC 4875 defines the following two approaches for handling the P2MP LSP re-merge:

- First, the node detecting the re-merge allows the re-merge case to persist, but data from all but one incoming interface is dropped at the re-merge node. This works by default without any configuration.
- Second, the re-merge node initiates the pruning of the re-merge sub LSPs through signaling.

On Juniper Networks MX Series routers, the first approach (as defined by RFC 4875) works by default. The second approach can be implemented by one of the following CLI configuration statements depending upon where the Juniper Networks MX Series routers are placed (ingress node or transit node) in the P2MP RSVP MPLS network:

- **no-re-merge**—This CLI configuration statement when enabled at the ingress (headend) router avoids the path computation of P2MP sub LSPs which creates a re-merge condition. When this CLI configuration statement is configured at the ingress, then configuring the **no-p2mp-re-merge** CLI configuration statement at the transit router is not required.
- **no-p2mp-re-merge**—This CLI configuration statement when enabled at the transit router changes the default behavior of allowing the P2MP sub LSP sessions re-merge to rejecting the re-merge. This CLI configuration statement is primarily required when the ingress (headend router) is not a Juniper Networks MX Series router.
- **single-abr**—This command when enabled reduces re-merge condition beyond the inter-area, or inter-domain, or inter-AS RSVP P2MP LSPs.

The following topology explains the re-merge behavior in a P2MP LSP network:

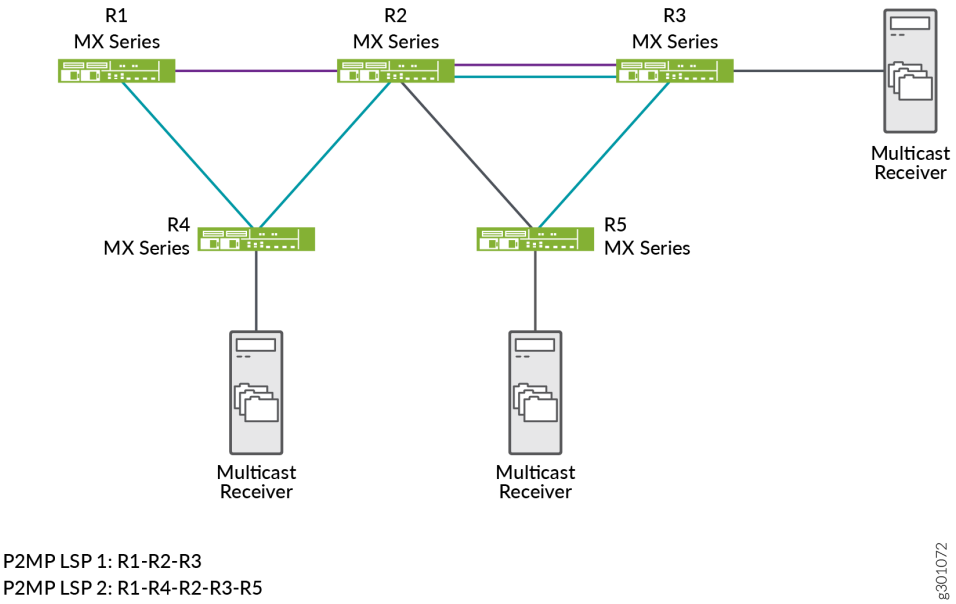

In this topology, R1 acts as an ingress (headend) router and R2 as the transit (re-merge node) router. There are two sub LSPs sessions created in this network, LSP 1 and LSP 2. LSP 1 is a session established between R1, R2, and R3 devices. LSP 2 is a session established between R1, R4, R2, R3, and R5 devices. By default, the transit router allows the re-merge to happen from both the sub LSPs and drops one of the sub LSP branch traffic at the re-merge node. You can control this re-merge behavior by enabling the **no-re-merge** CLI configuration statement at the ingress router, or the **no-p2mp-re-merge** CLI configuration statement at the transit router.

If you enable the **no-re-merge** CLI configuration statement at the ingress router (R1), only one of the two sub LSP session is established. For example, if LSP 1 (R1-R2-R3) session is established first, then the other sub LSP session (LSP 2) will not be established.

If you enable the **no-p2mp-re-merge** CLI configuration statement at the transit router (R2), the transit router rejects the re-merge of one of the sub LSPs and sends a path error message to the ingress router (R1) preventing the ingress router to not create the second P2MP LSP re-merge branch. You can use the **show rsvp statistics** CLI command to view the path error message.

Modify the Default P2MP LSP Re-merge Behavior

You can modify the default re-merge behavior either at the ingress (headend) node, or at the transit node in a P2MP RSVP MPLS network.

On the ingress (headend router), disable the default re-merge behavior so that the ingress router does not do the path computation of sub LSPs which creates the re-merge condition. The default behavior allows the path computation of sub LSPs.

```
[edit protocols]
user@host#set mpls p2mp-lsp no-re-merge
```

On the transit router, disable the default re-merge behavior so that the transit router rejects the re-merge of sub LSPs.

```
[edit protocols]
user@host#set rsvp no-p2mp-re-merge
```

For inter-area, or inter-domain, or inter-AS RSVP P2MP LSPs, use the **single-abr** CLI configuration statement at the ingress (headend router) so that all the P2MP sub LSP prefers to select the same exit router (ABR or ASBR) thereby reducing the re-merge condition.

```
[edit protocols]
user@host#set mpls p2mp-lsp single-abr
```

WHAT'S NEXT

For more information on P2MP LSPs, see Point-to-Multipoint LSP Configuration.