JUNIPER NETWORKS | Engineering Simplicity

Junos® OS

Logical Systems and Tenant Systems User Guide for Security Devices

Published
2023-03-13

JUNOS

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

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, 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.

Junos® OS Logical Systems and Tenant Systems User Guide for Security Devices
Copyright © 2023 Juniper Networks, Inc. All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at 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

About This Guide | xxiv

1 Overview

Logical Systems and Tenant Systems Overview | 2

2 Logical Systems

Logical Systems Overview | 5

- Understanding Logical Systems for SRX Series Services Gateways | 5
- Features and Limitations of Logical Systems | 8
- Understanding Licenses for Logical Systems and Tenant Systems on SRX Series Devices | 10
- Understanding the Interconnect Logical System and Logical Tunnel Interfaces | 11
- Understanding Packet Flow in Logical Systems for SRX Series Devices | 12
- Logical Systems and Tenant Systems support for VSRX and VSRX 3.0 Instances | 21

Primary Logical Systems Overview | 23

- Understanding the Primary Logical Systems and the Primary Administrator Role | 23
- SRX Series Logical Systems Primary Administrator Configuration Tasks Overview | 24
- Example: Configuring Multiple VPLS Switches and LT Interfaces for Logical Systems | 27
  - Requirements | 28
  - Overview | 28
  - Configuration | 30
  - Verification | 48

User Logical Systems Overview | 51

- User Logical Systems Configuration Overview | 51
- Understanding User Logical Systems and the User Logical System Administrator Role | 53

Setting Up a Logical System | 55

- Example: Configuring Root Password for Logical Systems | 55
  - Requirements | 55

Overview | 55

Configuration | 56

Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System | 57

Requirements | 57

Overview | 57

Configuration | 59

Verification | 67

Security Profiles for Logical Systems | 70

Understanding Logical Systems Security Profiles (Primary Administrators Only) | 71

Example: Configuring Logical Systems Security Profiles (Primary Administrators Only) | 77

Requirements | 78

Overview | 78

Configuration | 78

Verification | 88

Example: Configuring User Logical Systems Security Profiles | 89

Requirements | 90

Overview | 90

Configuration | 92

Verification | 95

Example: Configuring Security log stream for Logical Systems | 96

Requirements | 97

Overview | 97

Configuration | 97

Verification | 98

CPU Allocation for Logical Systems | 103

Understanding CPU Allocation and Control | 103

Example: Configuring CPU Utilization (Primary Administrators Only) | 108

Requirements | 108

Overview | 108

Configuration | 109

Verification | 111

Routing and Interfaces for Primary Logical Systems | 113

Understanding Logical Systems Interfaces and Routing Instances | **113**

Example: Configuring Interfaces, Routing Instances, and Static Routes for the Primary and Interconnect Logical Systems and Logical Tunnel Interfaces for the User Logical Systems (Primary Administrators Only) | **114**

Requirements | **115**

Overview | **115**

Configuration | **118**

Verification | **126**

Example: Configuring OSPF Routing Protocol for the Primary Logical Systems | **127**

Requirements | **127**

Overview | **127**

Configuration | **128**

Verification | **130**

Routing, Interfaces, and NAT for User Logical Systems | 132

Understanding Logical Systems Network Address Translation | **133**

Example: Configuring Network Address Translation for a User Logical Systems | **134**

Requirements | **134**

Overview | **135**

Configuration | **135**

Verification | **138**

Example: Configuring Interfaces and Routing Instances for a User Logical Systems | **139**

Requirements | **139**

Overview | **139**

Configuration | **140**

Example: Configuring OSPF Routing Protocol for a User Logical Systems | **143**

Requirements | **143**

Overview | **143**

Configuration | **144**

Verification | **147**

Security Zones in Logical Systems | 149

Understanding Logical Systems Zones | **149**

Example: Configuring User Logical Systems | **150**

Requirements | **151**

Overview | **151**

Configuration | **155**

Verification | **167**

Example: Configuring Security Zones for a User Logical Systems | **167**

Requirements | **168**

Overview | **168**

Configuration | **169**

User Authentication for Logical Systems | 173

Example: Configuring Access Profiles (Primary Administrators Only) | **173**

Requirements | **173**

Overview | **174**

Configuration | **174**

Example: Configuring Security Features for the Primary Logical Systems | **176**

Requirements | **177**

Overview | **177**

Configuration | **179**

Verification | **184**

Understanding Logical System Firewall Authentication | **185**

Example: Configuring Firewall Authentication for a User Logical System | **187**

Requirements | **187**

Overview | **187**

Configuration | **188**

Verification | **192**

Understanding Integrated User Firewall support in a Logical System | **193**

Example: Configuring Integrated User Firewall Identification Management for a User Logical System | **194**

Requirements | **195**

Overview | **196**

Configuration | **196**

Verification | **202**

Example: Configure Integrated User Firewall in Customized Model for Logical System | 205

Requirements | 205

Overview | 206

Configuration | 206

Verification | 210

Security Policies for Logical Systems | 213

Understanding Logical Systems Security Policies | 213

Example: Configuring Security Policies in a User Logical Systems | 216

Requirements | 216

Overview | 217

Configuration | 218

Verification | 220

Configuring Dynamic Address for Logical Systems | 221

Screen Options for User Logical Systems | 223

Understanding Logical Systems Screen Options | 223

Example: Configuring Screen Options for a User Logical Systems | 224

Requirements | 224

Overview | 224

Configuration | 225

Secure Wire for Logical Systems | 227

Secure Wire for Logical Systems Overview | 227

Example: Configure Secure Wire for User Logical Systems | 229

Requirements | 229

Overview | 230

Configuration | 230

Verification | 231

VPNs in Logical Systems | 233

Understanding Route-Based VPN Tunnels in Logical Systems | 233

Example: Configuring IKE and IPsec SAs for a VPN Tunnel (Primary Administrators Only) | 235

Requirements | 235

Overview | 235

Configuration | **238**

Verification | **242**

Example: Configuring a Route-Based VPN Tunnel in a User Logical Systems | **245**

Requirements | **245**

Overview | **245**

Configuration | **246**

Verification | **249**

UTM for Logical Systems | 250

Understanding UTM Features in Logical Systems | **251**

Example: Configuring UTM for the Primary Logical System | **252**

Requirements | **252**

Overview | **252**

Configuration | **253**

Verification | **258**

Example: Configuring UTM for a User Logical System | **262**

Requirements | **262**

Overview | **263**

Configuration | **264**

Verification | **269**

IDP for Logical Systems | 273

IDP in Logical Systems Overview | **274**

Understanding IDP Features in Logical Systems | **276**

Example: Configuring an IDP Policy for the Primary Logical Systems | **280**

Requirements | **280**

Overview | **280**

Configuration | **282**

Verification | **287**

Example: Configuring and Assigning a Predefined IDP Policy for a User Logical System | **288**

Requirements | **289**

Overview | **289**

Configuration | **289**

Verification | **291**

Example: Enabling IDP in a User Logical System Security Policy | **292**

Requirements | **292**

Overview | **292**

Configuration | **293**

Verification | **295**

Example: Configuring an IDP Policy for a User Logical System | **296**

Requirements | **296**

Overview | **296**

Configuration | **297**

Verification | **303**

ALG for Logical Systems | 304

Understanding Application Layer Gateway (ALG) in Logical Systems | **305**

Enabling and Disabling ALG for Logical System | **306**

Example: Enabling FTP ALG in a Logical System | **311**

Requirements | **311**

Overview | **311**

Configuration | **312**

Verification | **318**

DHCP for Logical Systems | 323

Understanding DHCP Support for Logical Systems | **323**

Minimum DHCPv6 Relay Agent Configuration for Logical Systems | **323**

Example: Configuring the DHCPv6 Client for Logical Systems | **325**

Requirements | **325**

Overview | **326**

Configuration | **326**

Verification | **330**

Example: Configuring the DHCPv6 Server Options for Logical Systems | **333**

Requirements | **333**

Overview | **334**

Configuration | **334**

Verification | **338**

Application Security in Logical Systems | 338

Understanding Logical Systems Application Identification Services | **339**

Understanding Logical Systems Application Firewall Services | **341**

Example: Configuring Application Firewall Services for a Primary Logical Systems | **342**

- Requirements | **342**
- Overview | **343**
- Configuration | **343**
- Verification | **346**

Understanding Logical Systems Application Tracking Services | **348**

Example: Configuring Application Firewall Services for a User Logical System | **349**

- Requirements | **349**
- Overview | **350**
- Configuration | **350**
- Verification | **353**

Example: Configuring AppTrack for a User Logical Systems | **355**

- Requirements | **355**
- Overview | **355**
- Configuration | **356**
- Verification | **358**

IPv6 for Logical Systems | 359

IPv6 Addresses in Logical Systems Overview | **360**

Understanding IPv6 Dual-Stack Lite in Logical Systems | **361**

Example: Configuring IPv6 for the Primary, Interconnect, and User Logical Systems (Primary Administrators Only) | **362**

- Requirements | **362**
- Overview | **363**
- Configuration | **365**
- Verification | **372**

Example: Configuring IPv6 Zones for a User Logical Systems | **373**

- Requirements | **373**
- Overview | **373**

Configuration | 374

Example: Configuring IPv6 Security Policies for a User Logical Systems | 378

Requirements | 378

Overview | 379

Configuration | 380

Verification | 382

Example: Configuring IPv6 Dual-Stack Lite for a User Logical Systems | 383

Requirements | 383

Overview | 384

Configuration | 384

Verification | 385

SSL Proxy for Logical Systems | 387

Understanding SSL Forward and Reverse Proxy for Logical Systems | 387

Example: Configuring SSL Forward and Reverse Proxy for Logical Systems | 388

Requirements | 388

Overview | 388

Configuration | 388

Verification | 392

ICAP Redirects for Logical Systems | 393

ICAP Redirect Support for Logical Systems | 393

Example: Configuring ICAP Redirect Service on SRX Devices | 395

Requirements | 395

Overview | 395

Configuration | 396

Verification | 400

AppQoS for Logical Systems | 402

Application Quality of Service Support for Logical Systems Overview | 402

Example: Configure Application Quality of Service for Logical Systems | 403

Requirements | 403

Overview | 404

Configuration | 404

Verification | 408

Logical Systems in a Chassis Cluster | 410

Understanding Logical Systems in the Context of Chassis Cluster | **410**

Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (Primary Administrators Only) | **411**

Requirements | **411**

Overview | **412**

Configuration | **415**

Verification | **447**

Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (IPv6) (Primary Administrators Only) | **456**

Requirements | **456**

Overview | **457**

Configuration | **460**

Verification | **492**

Flow Trace for Logical Systems | 501

Flow Trace Support for Logical Systems Overview | **502**

Configure Flow Trace Support for Logical Systems | **502**

Example: Deleting a Logical System | 503

Requirements | **504**

Overview | **504**

Configuration | **505**

Verification | **508**

Troubleshooting Logical Systems | 508

Understanding Security Logs and Logical Systems | **509**

Configuring On-Box Reporting for logical Systems | **511**

Example: Configure Security Log for Logical Systems | **512**

Requirements | **512**

Overview | **512**

Configuration | **513**

Verification | **516**

Configuring On-Box Binary Security Log Files for Logical System | 517

Configuring Off-Box Binary Security Log Files for Logical System | 519

Understanding Data Path Debugging for Logical Systems | 520

Performing Tracing for Logical Systems (Primary Administrators Only) | 521

Troubleshooting DNS Name Resolution in Logical System Security Policies (Primary Administrators Only) | 527

3 Tenant Systems

Tenant Systems Overview | 530

Understanding Tenant Systems | 530

Tenant System Configuration Overview | 537

Configuring a Routing Instance for a Tenant System | 539

Understanding Routing and Interfaces for Tenant Systems | 541

Overview: Configuring Routing and Interfaces for Tenant Systems | 541

Understanding Tenant System Security Profiles (Primary Administrators Only) | 548

Example: Creating Tenant Systems, Tenant System Administrators, and an Interconnect VPLS Switch | 554

Requirements | 554

Overview | 555

Full SRX Quick Configuration | 556

Verification | 569

Security Zones for Tenant Systems | 572

Understanding Zones for Tenant Systems | 573

Example: Configuring Zones in the Tenant System | 574

Requirements | 574

Overview | 574

Configuration | 575

Verification | 577

Flow for Tenant Systems | 578

Session Creation for Devices Running Tenant Systems | 579

Configuring Logical Systems and Tenant Systems Interconnect with Multiple VPLS Switches | 585

Requirements | **585**

Overview | **585**

Configuration | **586**

Verification | **594**

Configuring tenant systems Interconnect with Logical Tunnel Interface point-to-point connection | **596**

Requirements | **596**

Overview | **596**

Configuration | **597**

Verification | **605**

Configuring Logical System and Tenant System Interconnect with a Logical Tunnel Interface point-to-point connection | **606**

Requirements | **606**

Overview | **606**

Configuration | **607**

Verification | **612**

Flow Trace for Tenant Systems | 614

Flow Trace Support for Tenant Systems Overview | **615**

Configure Flow Trace Support for Tenant Systems | **615**

Firewall Authentication for Tenant Systems | 617

Understanding Tenant System Firewall Authentication | **617**

Configuring Firewall Authentication for a Tenant System | **620**

Requirements | **620**

Overview | **620**

Configuration | **622**

Verification | **633**

Understanding Integrated User Firewall Support in a Tenant System | **635**

Example: Configuring Integrated User Firewall Identification Management for a Tenant System | **636**

Requirements | **637**

Overview | **637**

Configuration | **637**

Verification | **644**

Example: Configure Integrated User Firewall in Customized Model for Tenant System | **646**

Requirements | **647**

Overview | **647**

Configuration | **647**

Verification | **651**

Security Policies for Tenant Systems | 654

Understanding Security Policies for Tenant Systems | **654**

Example: Configuring Security Policies in the Tenant System | **656**

Requirements | **656**

Overview | **657**

Configuration | **658**

Verification | **660**

Configuring Dynamic Address for Tenant Systems | **662**

Screen Options for Tenant Systems | 664

Understanding Tenant System Screen Options | **664**

Example: Configuring Screen Options for a Tenant System | **665**

Requirements | **665**

Overview | **665**

Configuration | **666**

Verification | **670**

NAT for Tenant Systems | 672

Understanding Network Address Translation for Tenant systems | **672**

Example: Configuring Network Address Translation for the Tenant Systems | **673**

Requirements | **673**

Overview | **674**

Configuration | **675**

Verification | **679**

UTM for Tenant Systems | 682

Understanding UTM Features in Tenant Systems | **683**

Example: Configuring UTM for the Tenant System | **684**

Requirements | **684**

Overview | **684**

Configuration | **685**

Verification | **689**

IDP for Tenant Systems | 690

Understanding IDP for Tenant Systems | **690**

Understanding IDP Features in Tenant Systems | **692**

Example: Configuring IDP Policies and Attacks for Tenant Systems | **694**

Requirements | **694**

Overview | **695**

Configuration | **695**

Verification | **709**

ALG for Tenant Systems | 713

Understanding ALG Support for Tenant System | **713**

Enabling and Disabling ALG for Tenant System | **714**

Example: Configuring ALG in Tenant System | **719**

Requirements | **720**

Overview | **720**

Configuration | **720**

Verification | **725**

DHCP for Tenant Systems | 727

Understanding DHCP support for Tenant Systems | **727**

Minimum DHCPv6 Relay Agent Configuration for Tenant Systems | **727**

Example: Configuring a DHCPv6 Client for Tenant Systems | **729**

Requirements | **729**

Overview | **729**

Configuration | **730**

Verification | **734**

Security Log for Tenant Systems | 737

Understanding of Security Log for Tenant Systems | **737**

Example: Configure Security Log for Tenant Systems | **739**

Requirements | 739
Overview | 739
Configuration | 740
Verification | 743

Understanding On-Box Reporting for Tenant Systems | 744

Configuring On-Box Reporting for Tenant Systems | 745

Understanding On-Box and Off-Box Logging for Tenant System | 746

Configuring On-Box Binary Security Log Files for Tenant System | 747

Configuring Off-Box Binary Security Log Files for Tenant System | 748

AppQoS for Tenant Systems | 750

Application Quality of Service for Tenant Systems Overview | 750

Example: Configure Application Quality of Service for Tenant Systems | 751
Requirements | 751
Overview | 752
Configuration | 752
Verification | 756

Application Security for Tenant Systems | 758

Application Identification Services for Tenant Systems Overview | 758

4 **Configuration Statements**

address-book (System) | 763

address-name | 765

anti-spam | 767

anti-virus | 769

auth-entry | 773

confidentiality | 775

content-filtering (Logical System Security Feature Profile) | 776

cpu | 779

dslite-softwire-initiator | 781

dynamic-address | 783

firewall-authentication (tenants) | 785

web-authentication | 787

pass-through | 789

flow-gate | 791

flow-session | 793

idp (logical-systems) | 796

idp-policy | 797

log (Security) | 799

log (Logical Systems and Tenant Systems) | 804

logical-system (System Security Profile) | 808

logical-domain-identity-management | 810

logical-systems (All) | 813

nat | 815

nat-cone-binding | 820

nat-destination-pool | 823

nat-destination-rule | 825

nat-interface-port-ol (System) | 827

nat-nopat-address | 829

nat-pat-address | 831

nat-pat-portnum | 833

nat-port-ol-ipnumber | 835

nat-rule-referenced-prefix (System) | 837

nat-source-pool | 839

nat-source-rule | 841

nat-static-rule | 843

policy (System Security Profile) | 845

policy-with-count | 848

profile | 850

protocols (Tenant Systems) | 851

purging | 853

root-authentication | 854

root-logical-system | 857

root-streaming | 858

secure-wire (System Security Profile) | 861

scheduler (System Security Profile) | 862

screen | 865

security-profile | 871

security-profile-resources | 876

stream (Logical Systems and Tenant Systems) | 878

softwires | 881

url | 883

web-filtering (Logical System Security Feature Profile) | 884

zone (System Security Profile) | 889

5 Operational Commands

clear class-of-service application-traffic-control counter | 896

clear class-of-service application-traffic-control rate-limiters | 898

clear class-of-service application-traffic-control statistics rule | 900

clear security application-firewall rule-set statistics logical-system | 902

clear security dns-cache | 904

clear security firewall-authentication users | 906

clear security firewall-authentication history | 909

clear security idp attack table | 912

clear security idp counters ips | 913

clear security idp counters pdf-decoder | 915

clear security idp counters ssl-inspection | 917

clear security idp counters memory | 919

clear security idp counters memory | 921

clear security idp counters tcp-reassembler | 922

clear security idp counters application-identification | 924

clear security idp counters action | 926

clear security idp counters dfa | 928

clear security idp counters flow | 929

clear security idp counters log | 931

clear security idp counters http-decoder | 933

clear security idp counters packet-log | 934

clear security idp counters packet | 936

clear security idp counters policy-manager | 938

clear security flow session tenant | 940

clear services user-identification logical-domain-identity-management counters | 942

request security datapath-debug capture start | 944

request security datapath-debug capture stop | 945

set chassis cluster cluster-id node node-number reboot | 947

show chassis cluster status | 949

show class-of-service application-traffic-control rate-limiters | 954

show log | 966

show route tenant | 974

show security application-firewall rule-set | 977

show security application-firewall rule-set logical-system | 982

show security application-tracking counters | 987

show security alg status logical-system | 989

show security datapath-debug capture | 994

show security datapath-debug counter | 996

show security dns-cache | 998

show security dynamic-address | 1001

show security firewall-authentication history | 1011

show security firewall-authentication users | 1015

show security flow session | 1019

show security flow session tenant | 1034

show security idp logical system | 1037

show security idp attack table | 1039

show security idp counters action | 1042

show security idp counters application-identification | 1046

show security idp counters memory | 1053

show security idp counters ssl-inspection | 1057

show security idp counters pdf-decoder | 1061

show security idp counters log | 1065

show security idp counters ips | 1073

show security idp counters dfa | 1081

show security idp counters flow | 1084

show security idp counters http-decoder | 1101

show security idp counters packet-log | 1105

show security idp counters packet | 1120

show security idp counters policy-manager | 1130

show security idp counters tcp-reassembler | 1133

show security idp logical-system policy-association | 1141

show security idp policies | 1143

show security idp policy-commit-status | 1145

show security idp policy-templates-list | 1148

show security idp security-package-version | 1150

show security ike security-associations | 1153

show security ipsec security-associations | 1171

show security log report | 1209

show security match-policies | 1211

show security nat destination rule | 1221

show security nat destination summary | 1226

show security nat source rule | 1230

show security nat source summary | 1237

show security nat static rule | 1245

show security policies | 1252

show security screen statistics | 1276

show services user-identification logical-domain-identity-management | 1292

show system security-profile | 1296

show system security-profile secure-wire | 1305

show system security-profile scheduler | 1310

show system security-profile security-log-stream-number detail | 1315

show system security-profile security-log-stream-number | 1318

show system security-profile security-log-stream-number summary | 1322

show security softwires | 1325

show security softwires map-e confidentiality status | 1328

show security softwires map-e domain | 1329

show security zones | 1332

About This Guide

Use this guide to configure logical systems and tenant Systems in Junos OS on the SRX Series devices to partition a single device into multiple domains to perform security and routing functions.

CHAPTER 1

Overview

Logical Systems and Tenant Systems Overview | 2

Logical Systems and Tenant Systems Overview

With the Junos operating system (Junos OS) on SRX Series device, you can partition a single security device into multiple logical devices that can perform independent tasks. Because logical systems perform a subset of the tasks once handled by the main device, logical systems offer an effective way to maximize the use of a single security platform.

A complex network design requires multiple layers of switches, routers, and security devices, which might lead to challenges in maintenance, configuration, and operation. To reduce such complexity, Juniper Networks supports logical systems. Logical systems perform a subset of the actions of the main device and have their own unique routing tables, interfaces, policies, and routing instances.

For SRX Series devices, you can partition a single device into following secure contexts:

- Logical systems
- Tenant systems

Each logical system has its own discrete administrative domain, logical interfaces, routing instances, security firewall and other security features. A tenant system provides logical partitioning of the SRX device into multiple domains similar to logical systems and provides high scalability.

2

CHAPTER

Logical Systems

Logical Systems Overview | 5

Primary Logical Systems Overview | 23

User Logical Systems Overview | 51

Setting Up a Logical System | 55

Security Profiles for Logical Systems | 70

CPU Allocation for Logical Systems | 103

Routing and Interfaces for Primary Logical Systems | 113

Routing, Interfaces, and NAT for User Logical Systems | 132

Security Zones in Logical Systems | 149

User Authentication for Logical Systems | 173

Security Policies for Logical Systems | 213

Screen Options for User Logical Systems | 223

Secure Wire for Logical Systems | 227

VPNs in Logical Systems | 233

UTM for Logical Systems | 250

IDP for Logical Systems | 273

ALG for Logical Systems | 304

DHCP for Logical Systems | 323

Application Security in Logical Systems | 338

IPv6 for Logical Systems | 359

SSL Proxy for Logical Systems | 387

ICAP Redirects for Logical Systems | 393

AppQoS for Logical Systems | 402

Logical Systems in a Chassis Cluster | 410

Flow Trace for Logical Systems | 501

Example: Deleting a Logical System | 503

Troubleshooting Logical Systems | 508

Logical Systems Overview

IN THIS SECTION

- Understanding Logical Systems for SRX Series Services Gateways | 5
- Features and Limitations of Logical Systems | 8
- Understanding Licenses for Logical Systems and Tenant Systems on SRX Series Devices | 10
- Understanding the Interconnect Logical System and Logical Tunnel Interfaces | 11
- Understanding Packet Flow in Logical Systems for SRX Series Devices | 12
- Logical Systems and Tenant Systems support for VSRX and VSRX 3.0 Instances | 21

Logical systems enable you to partition a single device into multiple secure contexts that perform independent tasks. For more information, see the following topics:

Understanding Logical Systems for SRX Series Services Gateways

Logical systems for SRX Series devices enable you to partition a single device into secure contexts. Each logical system has its own discrete administrative domain, logical interfaces, routing instances, security firewall and other security features. By transforming an SRX Series device into a multitenant logical systems device, you can give various departments, organizations, customers, and partners—depending on your environment—private use of portions of its resources and a private view of the device. Using logical systems, you can share system and underlying physical machine resources among discrete user logical systems and the primary logical system.

The top part of Figure 1 on page 6 shows the three main configuration components of a logical system. The lower part of the figure shows a single device with a primary logical system and discrete user logical systems.

Logical systems include both primary and user logical systems and their administrators. The roles and responsibilities of the primary administrator and those of a user logical system administrator differ greatly. This differentiation of privileges and responsibilities is considered role-based administration and control.

Figure 1: Understanding Logical Systems

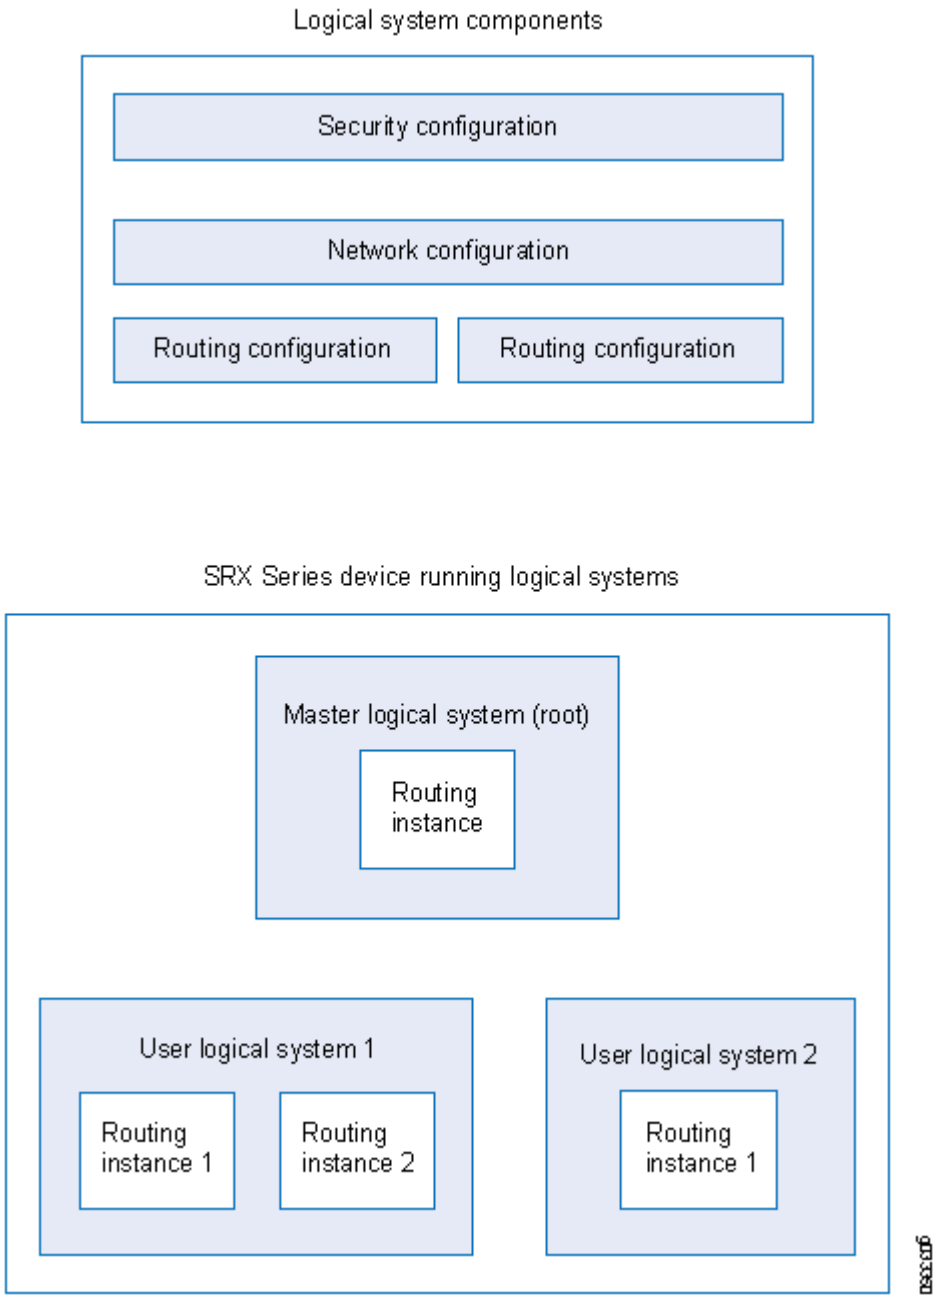

Logical systems on SRX Series devices offer many benefits, allowing you to:

- Curtail costs. Using logical systems, you can reduce the number of physical devices required for your company. Because you can consolidate services for various groups of users on a single device, you reduce both hardware costs and power expenditure.

- Create many logical systems on a single device and provision resources and services for them quickly. Because services are converged, it is easier for the primary, or root, administrator to manage a single device configured for logical systems than it is to manage many discrete devices.

You can deploy an SRX Series device running logical systems in many environments, in particular, in the enterprise and in the data center.

- In the enterprise, you can create and provision logical systems for various departments and groups.

  You can configure logical systems to enable communication among groups sharing the device. When you create logical systems for various departments on the same device, users can communicate with one another without traffic leaving the device if you have configured an interconnect logical system to serve as an internal switch. For example, members of the product design group, the marketing department, and the accounting department sharing an SRX Series Services Gateway running logical systems can communicate with one another just as they could if separate devices were deployed for their departments. You can configure logical systems to interconnect through *logical tunnel (lt-0/0/0)* internal interfaces. The lt-0/0/0 interfaces on the interconnect logical system connect to an lt-0/0/0 interface that you configure for each logical system. The interconnect logical system switches traffic between logical systems. The SRX Series device running logical systems provides for high, fast interaction among all logical systems created on the device when an interconnect logical system is used.

  Logical systems on the same device can also communicate with one another directly through ports on the device, as if they were separate devices. Although this method allows for direct connections between logical systems, it consumes more resources–you must configure interfaces and an external switch–and therefore it is more costly.

- In the data center, as a service provider, you can deploy an SRX Series device running logical systems to offer your customers secure and private user logical systems and discrete use of the device's resources.

  For example, one corporation might require 10 user logical systems and another might require 20. Because logical systems are secure, private, and self-contained, data belonging to one logical system cannot be viewed by administrators or users of other logical systems. That is, employees of one corporation cannot view the logical systems of another corporation.

- SRX4100 and SRX4200 devices support logical system in both transparent and route mode.

- SRX4600 device supports logical system in route mode only.

NOTE: To use the internal switch, which is optional, you must also configure an interconnect logical system. The interconnect logical system does not require an administrator.

NOTE: This feature requires a license. To understand more about SRX Series devices license, see Software Feature Licenses for SRX Series Devices. Please refer to the Juniper Licensing Guide for general information about License Management. Please refer to the product Data Sheets for details, or contact your Juniper Account Team or Juniper Partner.

SEE ALSO

Understanding the Primary Logical Systems and the Primary Administrator Role | 23

Understanding User Logical Systems and the User Logical System Administrator Role | 53

Features and Limitations of Logical Systems

This topic covers basic information about the features and limitations of logical systems.

- By default, logical systems deliver a primary logical system, which exists at the root level. You can purchase licenses for logical systems that you intend to create, with the total not exceeding 32 (1 primary logical system and 31 user logical systems).
- You can configure up to 32 security profiles, 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 of existing security profiles, commit the configuration, and then create the new security profiles and commit the configuration.
- You can configure one or more primary administrators to oversee administration of the device and the logical systems they configure.

  As primary administrator for an SRX Series Services Gateway running logical systems, you have root control over the device, its resources, and the logical systems that you create. You allocate security, networking, and routing resources to user logical systems. You can configure one logical system to serve as an interconnect logical system virtual private LAN service (VPLS) switch. The interconnect logical system, which is not mandatory, does not require security resources. However, if you configure an interconnect logical system, you must bind a dummy security profile to it. The primary administrator configures it and all lt-0/0/0 interfaces for it.

- A user logical system can have one or more administrators, referred to as user logical system administrators. The primary administrator creates login accounts for these administrators and assigns them to a user logical system. Currently, the primary administrator must configure all user logical system administrators. The first assigned user logical administrator cannot configure additional user logical system administrators for his or her logical system. As a user logical system administrator, you can configure the resources assigned to your user logical system, including logical interfaces assigned by the primary administrator, routing instances and their routes, and security components. You can display configuration information only for your logical system.
- A logical system can include more than one routing instance based on available system resources.
- You cannot configure *class of service* on lt-0/0/0 interfaces.
- Commit rollback is supported at the root level only.
- Quality-of-service (QoS) classification across interconnected logical systems does not work.
- The primary administrator can configure Application Layer Gateways (ALGs) at the root level. The configuration is inherited by all user logical systems. ALGs can also be configured discretely for user logical systems.
- The primary administrator can configure IDP policies at the root level and then apply an IDP policy to a user logical system.
- Only the primary administrator can create user accounts and login IDs for users for all logical systems. The primary administrator creates these user accounts at the root level and assigns them to the appropriate user logical systems.
- The same name cannot be used in two separate logical systems. For example, if logical-system1 includes a user with Bob configured as the username, then other logical systems on the device cannot include a user with the username Bob.
- Configuration for users for all logical systems and all user logical systems administrators must be performed at the root level by the primary administrator. A user logical system administrator cannot create other user logical system administrators or user accounts for their logical systems.
- Some of the scaling parameters are different for SRX1500 devices. For example, you can configure a maximum of 512 zones under a logical system.

SEE ALSO

Understanding Logical Systems for SRX Series Services Gateways | 5

Understanding the Primary Logical Systems and the Primary Administrator Role | 23

Understanding User Logical Systems and the User Logical System Administrator Role | 53

Understanding Licenses for Logical Systems and Tenant Systems on SRX Series Devices

This topic provides licensing information for SRX Series devices running logical systems and tenant systems.

Starting in Junos OS Release 18.3R1, an SRX Series device running logical systems or tenant systems includes three licenses by default. Starting in Junos OS Releases 20.3R3, 20.4R2, 21.1R2, and 21.2R1, you can use the features for logical systems and tenant systems without a license. One license for a primary logical system and the other two licenses for user-defined logical system or tenant system. The system does not allow you to configure additional logical systems or tenant systems if the number of logical systems and tenant systems exceeds the number of available licenses. In the earlier releases, the system allowed you to configure an additional logical system even if the number of logical systems exceeds the number of available licenses, but with a warning message of non-licensed logical-systems do not pass traffic. You can purchase licenses for additional logical systems and tenant systems that you intend to create. If you intend to configure an interconnect logical system or interconnect tenant system to use as a switch, it also requires separate licenses.

We enforce that you do not configure more logical systems or tenant systems than the number of licenses you have purchased. If the number of logical systems or tenant systems that you attempt to configure exceeds the number of licenses that you have purchased, then the system displays an error message similar to the following:

```
user@host# commit
error: 2 more multitenancy license(s) are needed!
error: configuration check-out failed
```

You can use the **show system license status all-logical-systems-tenants** or **show system license usage** commands to view the active logical systems and tenant systems on the device.

```
user@host> show system license status all-logical-systems-tenants
logical system name  license status
root-logical-system  enabled
LSYS2                enabled
LSYS0                enabled
LSYS11               enabled
LSYS12               enabled
LSYS23               enabled
TSYS1                enabled
```

```
TSYS2                enabled
TSYS3                enabled
```

```
user@host> show system license usage
                    Licenses  Licenses     Licenses   Expiry
  Feature name      used      installed    needed
  logical-system      9          11          0        2019-05-18
                                                      08:00:00 CST
```

When you use SRX Series devices running logical systems or tenant systems in a *chassis cluster*, you must purchase and install the same number of licenses for each node in the chassis cluster. Logical systems or tenant systems licenses pertain to a single chassis, or node, within a chassis cluster and not to the cluster collectively.

SEE ALSO

Logical Systems Overview | 5

Understanding User Logical Systems and the User Logical System Administrator Role | 53

Understanding the Interconnect Logical System and Logical Tunnel Interfaces

This topic covers the interconnect logical system that serves as an internal virtual private LAN service (VPLS) switch connecting one logical system on the device to another. The topic also explains how logical tunnel (lt-0/0/0) interfaces are used to connect logical systems through the interconnect logical system.

A device running logical systems can use an internal VPLS switch to pass traffic without it leaving the device. The interconnect logical system switches traffic across logical systems that use it. Although a virtual switch is used typically, it is not mandatory. If you choose to use a virtual switch, you must configure the interconnect logical system. There can be only one interconnect logical system on a device.

For communication between logical systems on the device to occur, you must configure an lt-0/0/0 interface on each logical system that will use the internal switch, and you must associate it with its peer lt-0/0/0 interface on the interconnect logical system, effectively creating a logical tunnel between them. You define a peer relationship at each end of the tunnel when you configure the logical system's lt-0/0/0 interfaces.

You might want all logical systems on the device to be able to communicate with one another without using an external switch. Alternatively, you might want some logical systems to connect across the internal switch but not all of them.

The interconnect logical system does not require security resources assigned to it through a security profile. However, you must assign a dummy security profile containing no resources to the interconnect logical system. Otherwise you will not be able to successfully commit the configuration for it.

WARNING: If you configure an lt-0/0/0 interface in any user logical system or the primary logical system and you do not configure an interconnect logical system containing a peer lt-0/0/0 interface for it, the commit will fail.

An SRX Series device running logical systems can be used in a *chassis cluster*. Each node has the same configuration, including the interconnect logical system.

When you use SRX Series devices running logical systems within a chassis cluster, you must purchase and install the same number of licenses for each node in the chassis cluster. Logical systems licenses pertain to a single chassis, or node, within a chassis cluster and not to the cluster collectively.

SEE ALSO

Example: Configuring Interfaces, Routing Instances, and Static Routes for the Primary and Interconnect Logical Systems and Logical Tunnel Interfaces for the User Logical Systems (Primary Administrators Only) | **114**

Understanding Logical Systems for SRX Series Services Gateways | **5**

Understanding Logical Systems in the Context of Chassis Cluster | **410**

Understanding Packet Flow in Logical Systems for SRX Series Devices

IN THIS SECTION

- Understanding Junos OS SRX Series Services Gateways Architecture | **14**
- Session Creation for Devices Running Logical Systems | **15**
- Understanding Flow on Logical Systems | **16**
- Understanding Packet Classification | **16**
- Handling Pass-Through Traffic for Logical Systems | **17**

- Handling Self-Traffic | **18**
- Understanding Session and Gate Limitation Control | **20**
- Understanding Sessions | **20**
- About Configuring Sessions | **20**

This topic explains how packets are processed in flow sessions on SRX Series devices running logical systems. It describes how an SRX Series device running logical systems handles pass-through traffic in a single logical system and between logical systems. It also covers self-traffic as self-initiated traffic within a logical system and self-traffic terminated on another logical system. Before addressing logical systems, the topic provides basic information about the SRX Series architecture in with respect to packet processing and sessions. Finally, it addresses sessions and how to change session characteristics.

The concepts explained in this example rely on the topology shown in Figure 2 on page 14.

Figure 2: Logical Systems, Their Virtual Routers, and Their Interfaces

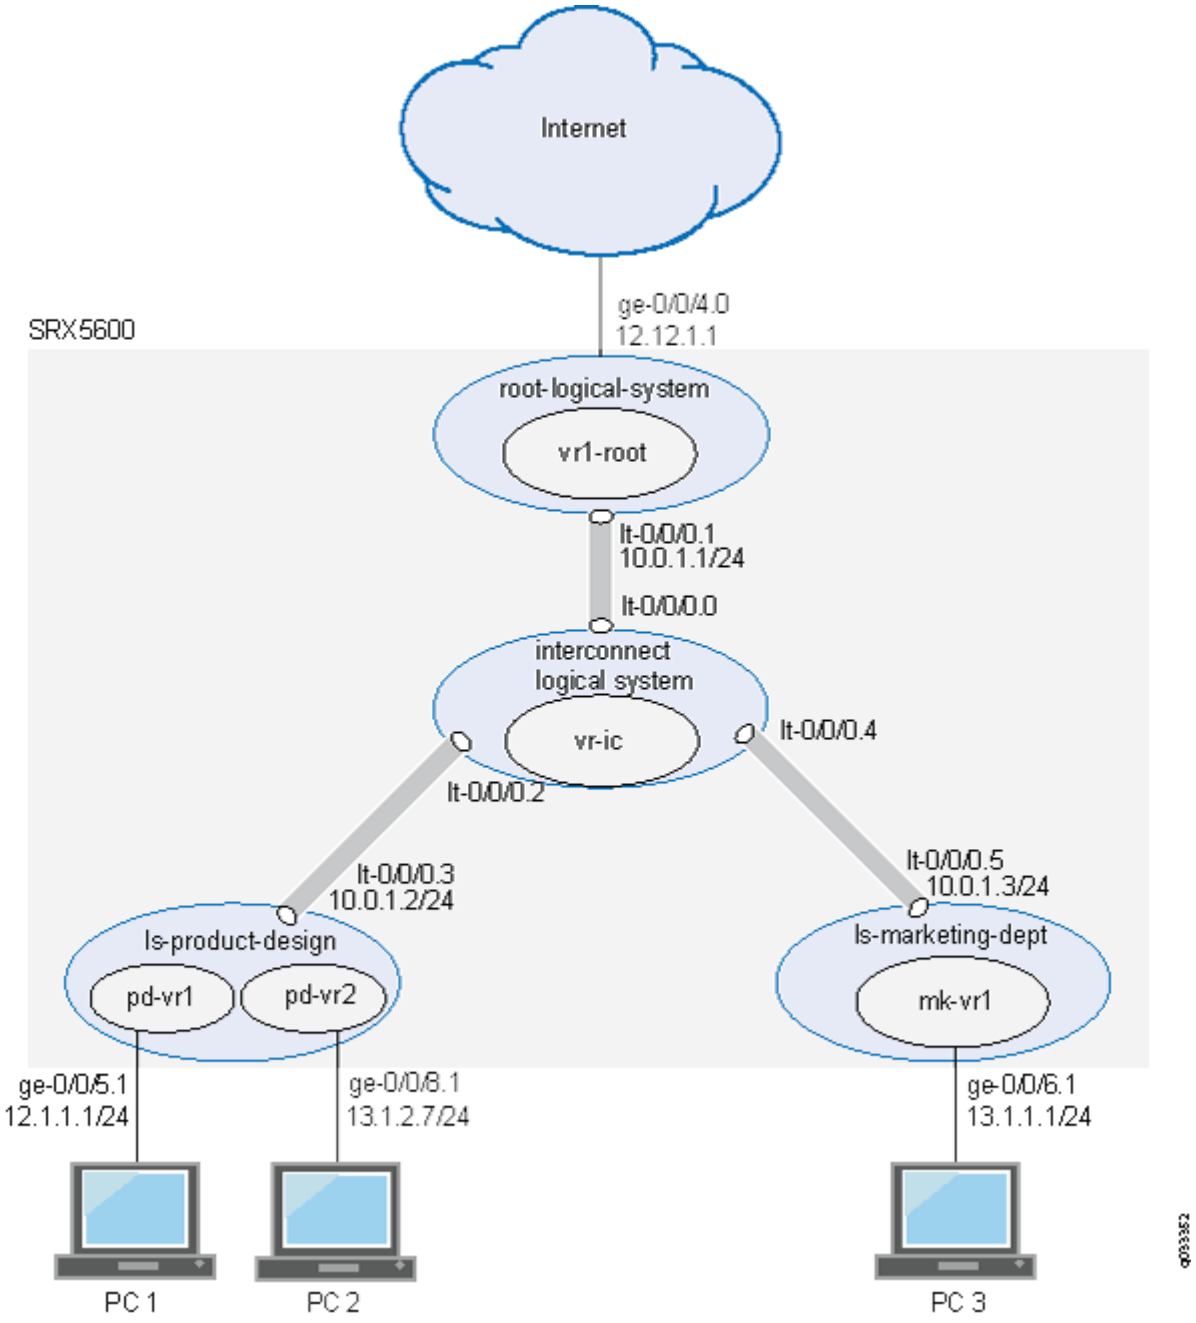

Understanding Junos OS SRX Series Services Gateways Architecture

Junos OS is a distributed parallel processing high throughput and high performance system. The distributed parallel processing architecture of the services gateways includes multiple processors to manage sessions and run security and other services processing. This architecture provides greater flexibility and allows for high throughput and fast performance.

The SRX5000 line devices include I/O cards (IOC) and Services Processing Cards (SPCs) that each contain processing units that process a packet as it traverses the device. A Network Processing Unit (NPU) runs on an IOC. An IOC has one or more NPUs. One or more Services Processing Units (SPUs) run on an SPC.

These processing units have different responsibilities. All flow-based services for a packet are executed on a single SPU. Otherwise, however, the lines are not clearly divided in regard to the kinds of services that run on these processors. (For details on flow-based processing, see *Understanding Traffic Processing on Security Devices*.)

For example:

- An NPU processes packets discretely. It performs sanity checks and applies some screens that are configured for the interface, such as denial-of-service (DoS) screens, to the packet.
- An SPU manages the session for the packet flow and applies security features and other services to the packet. It also applies packet-based stateless firewall filters, classifiers, and traffic shapers to the packet.
- The system uses one processor as a central point to take care of arbitration and allocation of resources and distribute sessions in an intelligent way. The central point assigns an SPU to be used for a particular session when the first packet of its flow is processed.

These discrete, cooperating parts of the system, including the central point, each store the information identifying whether a session exists for a stream of packets and the information against which a packet is matched to determine if it belongs to an existing session.

This architecture allows the device to distribute processing of all sessions across multiple SPUs. It also allows an NPU to determine if a session exists for a packet, to check the packet, and to apply screens to it. How a packet is handled depends on whether it is the first packet of a flow.

Flow-based packet processing treats related packets, or a stream of packets, in the same way. Packet treatment depends on characteristics that are established for the first packet of the packet stream when the flow session is established. Most packet processing occurs within a flow. For the distributed processing architecture of the services gateway, some packet-based processing, such as traffic shaping, occurs on the NPU. Some packet-based processing, such as application of classifiers to a packet, occurs on the SPU.

Configuration settings that determine the fate of a packet—such as the security policy that applies to it, Aplication Layer Gateway (ALG)s configured for it, if NAT should be applied to translate the packet's source and/or destination IP address—are assessed for the first packet of a flow.

Session Creation for Devices Running Logical Systems

Session establishment for SRX Series devices running logical systems differs in minor ways from that of SRX series devices not running logical systems. Despite the complexities that logical systems introduce,

traffic is handled in a manner similar to how it is handled on SRX Series devices not running logical systems. Flow-based packet processing, which is stateful, requires the creation of sessions. In considering flow based processing and session establishment for logical systems, it helps to think of each logical system on the device as a discrete device with respect to session establishment.

A session is created, based on routing and other classification information, to store information and allocate resources for a flow. Basically, a session is established when traffic enters a logical system interface, route lookup is performed to identify the next hop interface, and policy lookup is performed.

Optionally, logical systems enable you to configure an internal software switch. This virtual private LAN switch (VPLS) is implemented as an interconnect logical system. It enables both transit traffic and traffic terminated at a logical system to pass between logical systems. To enable traffic to pass between logical systems, logical tunnel (lt-0/0/0) interfaces across the interconnect logical system are used.

Communication between logical systems across the interconnect logical system requires establishment of two sessions: one for traffic that enters a logical system and exits its lt-0/0/0 interface, and one for traffic that enters the lt-0/0/0 interface of another logical system and either exits the device through one of its physical interface or is destined for it.

NOTE: Packet sequence occurs at the ingress and the egress interfaces. Packets traveling between logical systems might not be processed in the order in which they were received on the physical interface.

Understanding Flow on Logical Systems

To understand how traffic is handled for logical systems, it is helpful to consider each logical system as a discrete device.

NOTE: Traffic is processed for the primary logical system in the same way as it is for user logical systems on the device.

NOTE: On SRX1400, SRX3400, SRX3600, SRX5400, SRX5600, and SRX5800 Series devices, J-Flow version 5, version 8, and version 9 are not supported on logical systems.

Understanding Packet Classification

Packet classification is assessed the same way for SRX Series devices running with or without logical systems. Filters and class-of-service features are typically associated with an interface to influence

which packets are allowed to transit the system and to apply special actions to packets as needed. (Within a flow, some packet-based processing also takes place on an SPU.)

Packet classification is based on the incoming interface and performed at the ingress point. Traffic for a dedicated interface is classified to the logical system that contains that interface. Within the context of a flow, packet classification is based on both the physical interface and the *logical interface*.

Handling Pass-Through Traffic for Logical Systems

For SRX Series devices not running logical systems, pass-through traffic is traffic that enters and exits a device. You can think of pass-through traffic for logical systems similarly, but as having a larger dimension as a result of the nature of a multitenant device. For SRX Series devices running logical systems, pass-through traffic can exist within a logical system or between logical systems.

Pass-Through Traffic Within a Logical System

For pass-through traffic within a logical system, traffic comes in on an interface belonging to one of the logical system's virtual routing instances, and it is sent to another of its virtual routing instances. To exit the device, the traffic is sent out an interface belonging to the second virtual routing instance. The traffic does not transit between logical systems but rather enters and exits the device in a single logical system. Pass-through traffic within a logical system is transmitted according to the routing tables in each of its routing instances.

Consider how pass-through traffic is handled within a logical system given the topology shown in Figure 2 on page 14.

- When a packet arrives on interface ge-0/0/5, it is identified as belonging to the ls-product-design logical system.
- Because ge-0/0/5 belongs to the pd-vr1 routing instance, route lookup is performed in pd-vr1 with pd-vr2 identified as the next hop.
- A second route lookup is performed in pd-vr2 to identify the egress interface to use—in this case—ge-0/0/8.
- The packet is sent out ge-0/0/8 to the network.
- The security policy lookup is performed in ls-product-design, and one session is established.

Pass-Through Traffic Between Logical Systems

Pass-through traffic between logical systems is complicated by fact that each logical system has an ingress and an egress interface that the traffic must transit. It is as if traffic were coming into and going out from two devices.

Two sessions must be established for pass-through traffic between logical systems. (Note that policy lookup is performed in both logical systems).

- On the incoming logical system, one session is set up between the ingress interface (a physical interface) and its egress interface (an lt-0/0/0 interface).
- On the egress logical system, another session is set up between the ingress interface (the lt-0/0/0 interface of the second logical system) and its egress interface (a physical interface).

Consider how pass-through traffic is handled across logical systems in the topology shown in Figure 2 on page 14.

- A session is established in the incoming logical system.
  - When a packet arrives on interface ge-0/0/5, it is identified as belonging to the ls-product-design logical system.
  - Because ge-0/0/5 belongs to the pd-vr1 routing instance, route lookup is performed in pd-vr1.
  - As a result of the lookup, the egress interface for the packet is identified as lt-0/0/0.3 with the next hop identified as lt-0/0/0.5, which is the ingress interface in the ls-marketing-dept.
  - A session is established between ge-0/0/5 and lt-0/0/0.3.
- A session is established in the outgoing logical system.
  - The packet is injected into the flow again from lt-0/0/0.5, and the logical system context identified as ls-marketing-dept is derived from the interface.
  - Packet processing continues in the ls-marketing-dept logical system.
  - To identify the egress interface, route lookup for the packet is performed in the mk-vr1 routing instances.
  - The outgoing interface is identified as ge-0/0/6, and the packet is transmitted from the interface to the network.

Handling Self-Traffic

Self-traffic is traffic that originates in a logical system on the device and is either sent out to the network from that logical system or is terminated on another logical system on the device.

Self-Initiated Traffic

Self-initiated traffic is generated from a source logical system context and forwarded directly to the network from the logical system interface.

The following process occurs:

- When a packet is generated in a logical system, a process for handling the traffic is started in the logical system.
- Route lookup is performed to identify the egress interface, and a session is established.
- The logical system performs a policy lookup and processes the traffic accordingly.
- If required, a management session is set up.

Consider how self-initiated traffic is handled across logical systems given the topology shown in Figure 2 on page 14.

- A packet is generated in the ls-product-design logical system, and a process for handling the traffic is started in the logical system.
- Route lookup performed in pd-vr2 to identifies the egress interface as ge-0/0/8.
- A session is established.
- The packet is transmitted to the network from ge-0/0/8.

Traffic Terminated on a Logical System

When a packet enters the device on an interface belonging to a logical system and the packet is destined for another logical system on the device, the packet is forwarded between the logical systems in the same manner as is pass-through traffic. However, route lookup in the second logical system identifies the local egress interface as the packet destination. Consequently the packet is terminated on the second logical system as self-traffic.

- For terminated self-traffic, two policy lookups are performed, and two sessions are established.
    - On the incoming logical system, one session is set up between the ingress interface (a physical interface) and its egress interface (an lt-0/0/0 interface).
    - On the destination logical system, another session is set up between the ingress interface (the lt-0/0/0 interface of the second logical system) and the local interface.

Consider how terminated self-traffic is handled across logical systems in the topology shown in Figure 2 on page 14.

- A session is established in the incoming logical system.
    - When a packet arrives on interface ge-0/0/5, it is identified as belonging to the ls-product-design logical system.
    - Because ge-0/0/5 belongs to the pd-vr1 routing instance, route lookup is performed in pd-vr1.

    - As a result of the lookup, the egress interface for the packet is identified as lt-0/0/0.3 with the next hop identified as lt-0/0/0.5, the ingress interface in the ls-marketing-dept.
    - A session is established between ge-0/0/5 and lt-0/0/0.3.
- A management session is established in the destination logical system.
    - The packet is injected into the flow again from lt-0/0/0.5, and the logical system context identified as ls-marketing-dept is derived from the interface.
    - Packet processing continues in the ls-marketing-dept logical system.
    - Route lookup for the packet is performed in the mk-vr1 routing instance. The packet is terminated in the destination logical system as self-traffic.
    - A management session is established.

Understanding Session and Gate Limitation Control

The logical systems flow module provides session and gate limitation to ensure that these resources are shared fairly among the logical systems. Resources allocation and limitations for each logical system are specified in the security profile bound to the logical system.

- For session limiting, the system checks the first packet of a session against the maximum number of sessions configured for the logical system. If the maximum is reached, the device drops the packet and logs the event.
- For gate limiting, the device checks the first packet of a session against the maximum number of gates configured for the logical system. If the maximum number of gates for a logical system is reached, the device rejects the gate open request and logs the event.

Understanding Sessions

Sessions are created based on routing and other classification information to store information and allocate resources for a flow. You can change some characteristics of sessions, such as when a session is terminated. For example, you might want to ensure that a session table is never entirely full to protect against an attacker's attempt to flood the table and thereby prevent legitimate users from starting sessions.

About Configuring Sessions

Depending on the protocol and service, a session is programmed with a timeout value. For example, the default timeout for TCP is 1800 seconds. The default timeout for UDP is 60 seconds. When a flow is terminated, it is marked as invalid, and its timeout is reduced to 10 seconds. If no traffic uses the If no

traffic uses the session before the service timeout, the session is aged out and freed to a common resource pool for reuse.

You can affect the life of a session in the following ways:

- Age out sessions, based on how full the session table is.
- Set an explicit timeout for aging out TCP sessions.
- Configure a TCP session to be invalidated when it receives a TCP RST (reset) message.
- You can configure sessions to accommodate other systems as follows:
  - Disable TCP packet security checks.
  - Change the maximum segment size.

SEE ALSO

Understanding the Interconnect Logical System and Logical Tunnel Interfaces | 11

Understanding Logical Systems for SRX Series Services Gateways | 5

Logical Systems and Tenant Systems support for VSRX and VSRX 3.0 Instances

Starting in Junos OS Release 20.1R1, you can configure logical systems and tenant systems on vSRX and vSRX 3.0 instances.

Configuring each logical systems creates an extra routing protocol process (RPD), which is cpu and memory intensive. Starting in Junos OS Release 20.1R1-

- vSRX and vSRX3.0 instances with a memory capacity of less than 16GB support one root logical system.
- vSRX and vSRX3.0 instances with a memory capacity of 16GB or more supports logical systems but limits the logical systems to eight.

Table 1 on page 22 describes the number of logical systems and tenant systems supported on different memory capacities for vSRX and vSRX 3.0.

Table 1: Logical Systems and Tenant Systems supported on different Memory Capacities for VSRX and VSRX 3.0.

Type	4GB	8GB	16GB or more
Logical systems (includes the root logical system)	1	1	8
Tenant systems	0	0	42
Logical systems + Tenant systems (includes the root logical system).	1	1	50

NOTE: Only VSRX 3.0 instances support flexible security profiles based on the device memory. The maximum number of supported security profiles on VSRX 3.0 is related to its memory. For more information, see Security Profiles for Logical Systems.

Use the following command at the `[edit]` hierarchy level to ensure that there are at least two CPU's in the Routing Engine of vSRX and vSRX3.0 instances: `set security forwarding-options resource-manager cpu re 2`.

SEE ALSO

- Logical Systems Overview
- Tenant Systems Overview

Release History Table

Release	Description
20.3R3	Starting in Junos OS Releases 20.3R3, 20.4R2, 21.1R2, and 21.2R1, you can use the features for logical systems and tenant systems without a license.
18.3R1	Starting in Junos OS Release 18.3R1, an SRX Series device running logical systems or tenant systems includes three licenses by default.

Primary Logical Systems Overview

IN THIS SECTION

- Understanding the Primary Logical Systems and the Primary Administrator Role | 23
- SRX Series Logical Systems Primary Administrator Configuration Tasks Overview | 24
- Example: Configuring Multiple VPLS Switches and LT Interfaces for Logical Systems | 27

Primary logical systems can create a user logical system and configure the security resources of the user logical system. Primary logical systems assign the logical interfaces to the user logical systems. For more information, see the following topics:

Understanding the Primary Logical Systems and the Primary Administrator Role

When, as a primary administrator, you initialize an SRX Series device running logical systems, a primary logical system is created at the root level. You can log in to the device as root and change the root password.

By default, all system resources are assigned to the primary logical system, and the primary administrator allocates them to the user logical systems.

As primary administrator, you manage the device and all its logical systems. You also manage the primary logical system and configure its assigned resources. There can be more than one primary administrator managing a device running logical systems.

- The primary administrator's role and main responsibilities include:
  - Creating user logical systems and configuring their administrators. You can create one or more user logical system administrators for each user logical system.
  - Creating login accounts for users for all logical systems and assigning them to the appropriate logical systems.
  - Configuring an interconnect logical system if you want to allow communication between logical systems on the device. The interconnect logical system acts as an internal switch. It does not require an administrator.

To configure an interconnect logical system, you configure lt-0/0/0 interfaces between the interconnect logical system and each logical system. These peer interfaces effectively allow for establishment of tunnels.

- Configuring security profiles to provision portions of the system's security resources to user logical systems and the primary logical system.

  Only the primary administrator can create, change, and delete security profiles and bind them to logical systems.

  > **NOTE**: A user logical system administrator can configure interface, routing, and security resources allocated to his logical system.

- Creating logical interfaces to assign to user logical systems. (The user logical system administrator configures logical interfaces assigned to his logical system.)
- Viewing and managing user logical systems, as required, and deleting user logical systems. When a user logical system is deleted, its allocated reserved resources are released for use by other logical systems.
- Configuring IDP, AppTrack, application identification, and application firewall features. The primary administrator can also use trace and debug at the root level, and he can perform commit rollbacks. The primary administrator manages the primary logical system and configures all the features that a user logical system administrator can configure for his or her own logical systems including routing instances, static routes, dynamic routing protocols, zones, security policies, screens, and firewall authentication.

SEE ALSO

Understanding User Logical Systems and the User Logical System Administrator Role | 53

Understanding Logical Systems for SRX Series Services Gateways | 5

Example: Configuring Interfaces, Routing Instances, and Static Routes for the Primary and Interconnect Logical Systems and Logical Tunnel Interfaces for the User Logical Systems (Primary Administrators Only) | 114

SRX Series Logical Systems Primary Administrator Configuration Tasks Overview

This topic describes the primary administrator's tasks in the order in which they are performed.

An SRX Series device running logical systems is managed by a primary administrator. The primary administrator has the same capabilities as the root administrator of an SRX Series device not running logical systems. However, the primary administrator's role and responsibilities extend beyond those of other SRX Series device administrators because an SRX Series device running logical systems is partitioned into discrete logical systems, each with its own resources, configuration, and management concerns. The primary administrator is responsible for creating these user logical systems and provisioning them with resources.

For an overview of the primary administrator's role and responsibilities, see "Understanding the Primary Logical Systems and the Primary Administrator Role" on page 23.

As the primary administrator, you perform the following tasks to configure an SRX Series device running logical systems:

1. Configure a root password. Initially the primary administrator logs in to the device as the root user without needing to specify a password. After you log in to the device, you must define a root password for later use.

   See "Example: Configuring Root Password for Logical Systems" on page 55 for configuration information.

2. Create user logical systems and their administrators and users. Optionally, create an interconnect logical system.

   For each user logical system that you want to configure on the device, you must create a logical system, define one or more administrators for it, and add users to it.

   The primary administrator configures login accounts for user logical system administrators and users and associates them with the user logical system. A user logical system can have more than one administrator; the primary administrator must define and add all user logical system administrators and add them to their user logical systems.

   The primary administrator adds users to user logical systems on behalf of the user logical system administrator. For example, if you have created a user logical system for the product design department, you must create user accounts for the users who belong to that department and associate them with the user logical system. The user logical system administrator does not have the ability to do this. Rather, the user logical administrator tells you the user accounts that you must create and add for his logical system.

   - For configuration information, see "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System" on page 57
   - For information on user logical system administrators, see "Understanding User Logical Systems and the User Logical System Administrator Role" on page 53.
   - For information on the interconnect logical system, see "Understanding the Interconnect Logical System and Logical Tunnel Interfaces" on page 11.

3. Configure one or more security profiles. Security profiles assign security resources to logical systems. You can assign a single security profile to more than one logical system if you intend to allocate the same kinds and amounts of resources to them.

   - For configuration information, see "Example: Configuring Logical Systems Security Profiles (Primary Administrators Only)" on page 77.

   - For information on security profiles, see "Understanding Logical Systems Security Profiles (Primary Administrators Only)" on page 71.

4. Configure interfaces, routing instances, and static routes for logical systems, as appropriate.

   - If you plan to use an interconnect logical system, configure its logical tunnel interfaces and add them to its virtual routing instance.

   - Configure interfaces for the primary logical system. Optionally, create its logical tunnel interface to allow it to communicate with other logical systems on the device. Create a virtual routing instance for the primary logical system and add its interfaces and static routes to it. Also configure logical interfaces for user logical systems with VLAN tagging.

     NOTE: The primary administrator tells the user logical system administrators which interfaces are assigned to their logical systems. It is the user logical system administrator's responsibility to configure their interfaces.

   - Optionally, configure logical tunnel interfaces for any user logical systems that you want to allow to communicate with one another using the internal VPLS switch. VPLS is a virtual private network (VPN) technology. It allows point-to-point layer 2 tunnels connectivity.

     By creating a VPLS type routing-instance (RI), we define a VPLS switch. VPLS switch behaves like a L2 ethernet switch. We assign multiple LT IFLs to the VPLS switch. Each LT IFL have encapsulation ethernet-vpls and this behaves as L2 switch port. To connect to the VPLS switch, each logical system creates a LT IFL and assigns to a port of the VPLS switch.

     Starting with Junos OS Release 18.2R1, it is not required to define a dedicated interconnect logical system for including VPLS switch. For ease, VPLS switch is defined in root logical system. This approach is enabled by configuring multiple VPLS switches and LT IFLs per logical system.

     When one LT logical interface connects to a VPLS switch, the routing engine assigns VPLS switch unique MAC address from MAC address pool of the LT interface. This determines the number of LT IFLs that connect a VPLS switch.

   - For configuration information, see "Routing and Interfaces for Primary Logical Systems" on page 114.

   - For information about the interconnect logical system and logical tunnel (lt-0/0/0) interfaces, see "Understanding the Interconnect Logical System and Logical Tunnel Interfaces" on page 11.

5. Enable CPU utilization control and configure the CPU control target and reserved CPU quotas for logical systems. See "Example: Configuring CPU Utilization (Primary Administrators Only)" on page 108.
6. Optionally, configure dynamic routing protocols for the primary logical system. See "Example: Configuring OSPF Routing Protocol for the Primary Logical Systems" on page 127
7. Configure zones, security policies, and security features for the primary logical system. See "Example: Configuring Security Features for the Primary Logical Systems" on page 176.
8. Configure IDP for the primary logical system. See "Example: Configuring an IDP Policy for the Primary Logical Systems" on page 280.
9. Configure application firewall services on the primary logical system. See "Understanding Logical Systems Application Firewall Services" on page 341 and "Example: Configuring Application Firewall Services for a Primary Logical Systems" on page 342.
10. Configure a route-based VPN to secure traffic between a logical system and a remote site. See "Example: Configuring IKE and IPsec SAs for a VPN Tunnel (Primary Administrators Only)" on page 235.

SEE ALSO

Understanding Logical Systems for SRX Series Services Gateways | 5

Example: Configuring Multiple VPLS Switches and LT Interfaces for Logical Systems

IN THIS SECTION

- Requirements | 28
- Overview | 28
- Configuration | 30
- Verification | 48

This example shows how to interconnect multiple logical systems. This is achieved by configuring multiple logical systems with a Logical Tunnel (LT) interface point-to-point connection (Encapsulation Ethernet, Encapsulation Frame-Relay and Virtual Private LAN Service switch). More than one LT interface under a logical system and multiple VPLS switches are configured to pass the traffic without

leaving an SRX Series device. The frame-relay encapsulation adds data-link connection identifier (DLCI) information to the given frame.

Requirements

This example uses an SRX Series device running Junos OS with logical system.

Before you begin:

- Read the "SRX Series Logical Systems Primary Administrator Configuration Tasks Overview" on page 24 to understand how and where this procedure fits in the overall primary administrator configuration process.
- Read the "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System" on page 57
- Read the "Understanding the Interconnect Logical System and Logical Tunnel Interfaces" on page 11

Overview

In this example, we configure multiple LT interfaces and multiple VPLS switches under one logical system.

In this example, we also configure interconnect multiple logical systems with LT interface point-to point connection (Encapsulation Ethernet and Encapsulation Frame-Relay).

Figure 3 on page 28 shows the topology for interconnecting logical systems.

Figure 3: Configuring the interconnect logical systems

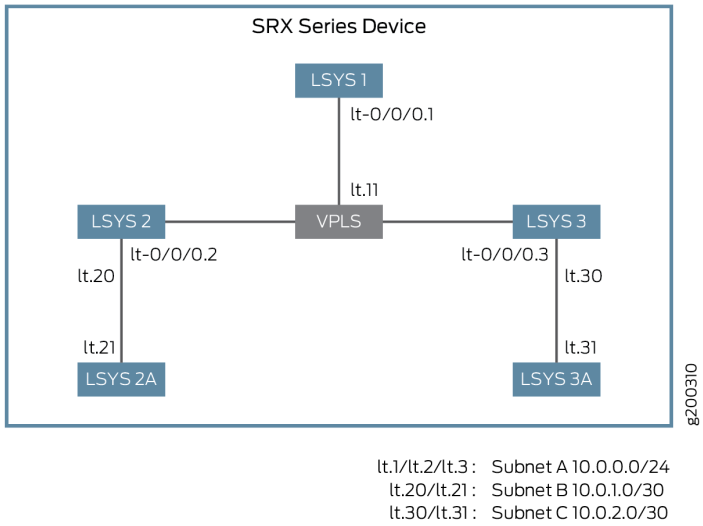

lt.1/lt.2/lt.3 : Subnet A 10.0.0.0/24
lt.20/lt.21 : Subnet B 10.0.1.0/30
lt.30/lt.31 : Subnet C 10.0.2.0/30

- For the interconnect logical system with LT interface point-to-point connection (encapsulation ethernet), the example configures logical tunnel interfaces lt-0/0/0. This example configures security-zone and assigns interfaces to the logical systems.

  The interconnect logical systems lt-0/0/0 interfaces are configured with Ethernet as the encapsulation type. The corresponding peer lt-0/0/0 interfaces in the logical systems are configured with Ethernet as the encapsulation type. A security profile is assigned to the logical systems.

- For the interconnect logical systems with LT interface point-to-point connection (encapsulation frame-relay), this example configures logical tunnel interfaces lt-0/0/0. This example configures security-zone and assigns interfaces to the logical systems.

  The interconnect logical systems lt-0/0/0 interfaces are configured with frame-relay as the encapsulation type. The corresponding peer lt-0/0/0 interfaces in the logical systems are configured with frame-relay as the encapsulation type. A security profile is assigned to the logical systems.

- For interconnect logical systems with multiple VPLS switches, this example configures logical tunnel interfaces lt-0/0/0 with ethernet-vpls as the encapsulation type. The corresponding peer lt-0/0/0 interfaces and security-profiles are assigned to the logical systems. The routing instance for the VPLS switch-1 and VPLS switch-2 are also assigned to the logical systems.

  Figure 4 on page 30 shows the topology for interconnect logical systems with VPLS switches.

Figure 4: Configuring the interconnect logical systems with VPLS switches

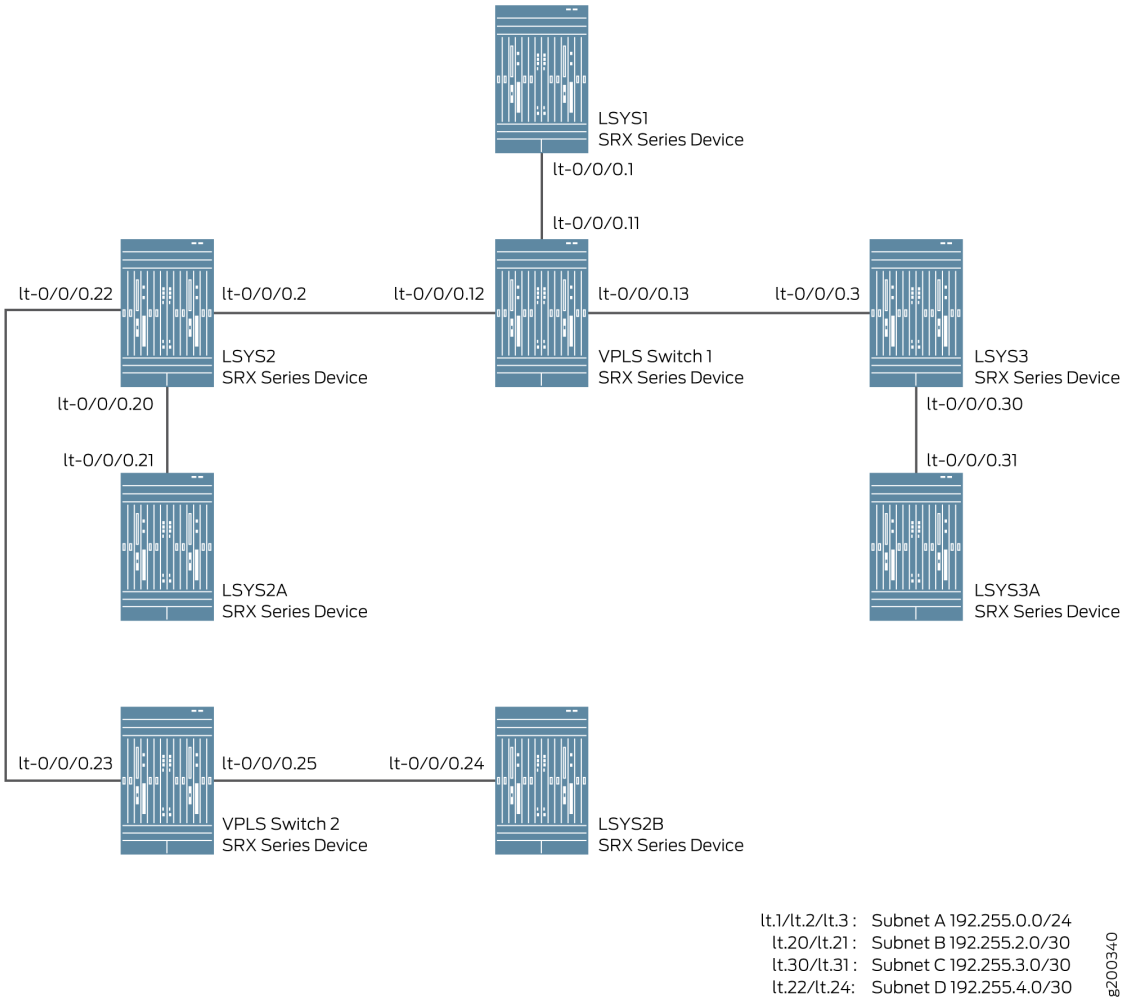

NOTE: Multiple LT interfaces can be configured within a logical system.

Configuration

IN THIS SECTION

- Configuring Logical Systems Interconnect with Logical Tunnel Interface point-to-point connection (Encapsulation Ethernet) | **31**

- Configuring Logical Systems Interconnect with Logical Tunnel Interface point-to-point connection (Encapsulation Frame-Relay) | 36
- Configuring Logical Systems Interconnect with Multiple VPLS Switches | 41

To configure interfaces for the logical system, perform these tasks:

Configuring Logical Systems Interconnect with Logical Tunnel Interface point-to-point connection (Encapsulation Ethernet)

CLI Quick Configuration

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

```
set system security-profile SP-user logical-system LSYS2
set logical-systems LSYS2 interfaces lt-0/0/0 unit 20 encapsulation ethernet
set logical-systems LSYS2 interfaces lt-0/0/0 unit 20 peer-unit 21
set logical-systems LSYS2 interfaces lt-0/0/0 unit 20 family inet address 192.255.2.1/30
set logical-systems LSYS2 security zones security-zone LT interfaces lt-0/0/0.20
set system security-profile SP-user logical-system LSYS2A
set logical-systems LSYS2A interfaces lt-0/0/0 unit 21 encapsulation ethernet
set logical-systems LSYS2A interfaces lt-0/0/0 unit 21 peer-unit 20
set logical-systems LSYS2A interfaces lt-0/0/0 unit 21 family inet address 192.255.2.2/30
set logical-systems LSYS2A security policies from-zone LT to-zone LT policy LT match source-
address any
set logical-systems LSYS2A security policies from-zone LT to-zone LT policy LT match destination-
address any
set logical-systems LSYS2A security policies from-zone LT to-zone LT policy LT match application
any
set logical-systems LSYS2A security policies from-zone LT to-zone LT policy LT then permit
set logical-systems LSYS2A security policies default-policy permit-all
set logical-systems LSYS2A security zones security-zone LT host-inbound-traffic system-services
all
set logical-systems LSYS2A security zones security-zone LT host-inbound-traffic protocols all
set logical-systems LSYS2A security zones security-zone LT interfaces lt-0/0/0.21
```

Step-by-Step Procedure

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

1. Define a security profile and assign to a logical system.

```
[edit]
user@host# set system security-profile SP-user logical-system LSYS2
```

2. Set the LT interface as encapsulation ethernet in the logical system.

```
[edit]
user@host# set logical-systems LSYS2 interfaces lt-0/0/0 unit 20 encapsulation ethernet
```

3. Configure a peer relationship for logical systems LSYS2.

```
[edit]
user@host# set logical-systems LSYS2 interfaces lt-0/0/0 unit 20 peer-unit 21
```

4. Specify the IP address for the LT interface.

```
[edit]
user@host# set logical-systems LSYS2 interfaces lt-0/0/0 unit 20 family inet address
192.255.2.1/30
```

5. Set the security zone for the LT interface.

```
[edit]
user@host# set logical-systems LSYS2 security zones security-zone LT interfaces lt-0/0/0.20
```

6. Define a security profile and assign to a logical system.

```
[edit]
user@host# set system security-profile SP-user logical-system LSYS2A
```

7. Set the LT interface as encapsulation ethernet in the logical system 2A.

```
[edit]
user@host# set logical-systems LSYS2A interfaces lt-0/0/0 unit 21 encapsulation ethernet
```

8. Configure a peer relationship for logical systems LSYS2A.

```
[edit]
user@host# set logical-systems LSYS2A interfaces lt-0/0/0 unit 21 peer-unit 20
```

9. Specify the IP address for the LT interface.

```
[edit]
user@host# set logical-systems LSYS2A interfaces lt-0/0/0 unit 21 family inet address
192.255.2.2/30
```

10. Configure a security policy that permits traffic from the LT zone to the LT policy LT zone.

```
[edit]
user@host# set logical-systems LSYS2A security policies from-zone LT to-zone LT policy LT
match source-address any
user@host# set logical-systems LSYS2A security policies from-zone LT to-zone LT policy LT
match destination-address any
user@host# set logical-systems LSYS2A security policies from-zone LT to-zone LT policy LT
match application any
user@host# set logical-systems LSYS2A security policies from-zone LT to-zone LT policy LT
then permit
```

11. Configure a security policy that permits traffic from default-policy.

```
[edit]
user@host# set logical-systems LSYS2A security policies default-policy permit-all
```

12. Configure security zones.

```
[edit]
user@host# set logical-systems LSYS2A security zones security-zone LT host-inbound-traffic
```

```
system-services all
user@host# set logical-systems LSYS2A security zones security-zone LT host-inbound-traffic
protocols all
user@host# set logical-systems LSYS2A security zones security-zone LT interfaces
lt-0/0/0.21
```

Results

- From configuration mode, confirm your configuration by entering the `show logical-systems LSYS2` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show logical-systems LSYS2
interfaces {
    lt-0/0/0 {
        unit 20 {
            encapsulation ethernet;
            peer-unit 21;
            family inet {
                address 192.255.2.1/30;
            }
        }
        unit 22 {
            encapsulation ethernet;
            peer-unit 23;
            family inet {
                address 192.255.4.1/30;
            }
        }
    }
}
security {
    zones {
        security-zone LT {
            interfaces {
                lt-0/0/0.22;
                lt-0/0/0.20;
            }
        }
```

```
    }
}
```

- From configuration mode, confirm your configuration by entering the `show logical-systems LSYS2A` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show logical-systems LSYS2A
interfaces {
    lt-0/0/0 {
        unit 21 {
            encapsulation ethernet;
            peer-unit 20;
            family inet {
                address 192.255.2.2/30;
            }
        }
    }
}
security {
    policies {
        from-zone LT to-zone LT {
            policy LT {
                match {
                    source-address any;
                    destination-address any;
                    application any;
                }
                then {
                    permit;
                }
            }
        }
        default-policy {
            permit-all;
        }
    }
    zones {
        security-zone LT {
            host-inbound-traffic {
                system-services {
```

```
                    all;
                }
                protocols {
                    all;
                }
            }
            interfaces {
                lt-0/0/0.21;
            }
        }
    }
}
```

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

Configuring Logical Systems Interconnect with Logical Tunnel Interface point-to-point connection (Encapsulation Frame-Relay)

CLI Quick Configuration

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

```
set system security-profile SP-user logical-system LSYS3A
set logical-systems LSYS3 interfaces lt-0/0/0 unit 30 encapsulation frame-relay
set logical-systems LSYS3 interfaces lt-0/0/0 unit 30 dlci 16
set logical-systems LSYS3 interfaces lt-0/0/0 unit 30 peer-unit 31
set logical-systems LSYS3 interfaces lt-0/0/0 unit 30 family inet address 192.255.3.1/30
set logical-systems LSYS3 security zones security-zone LT interfaces lt-0/0/0.30
set logical-systems LSYS3A interfaces lt-0/0/0 unit 31 encapsulation frame-relay
set logical-systems LSYS3A interfaces lt-0/0/0 unit 31 dlci 16
set logical-systems LSYS3A interfaces lt-0/0/0 unit 31 peer-unit 30
set logical-systems LSYS3A interfaces lt-0/0/0 unit 31 family inet address 192.255.3.2/30
set logical-systems LSYS3A security policies from-zone LT to-zone LT policy LT match source-
address any
set logical-systems LSYS3A security policies from-zone LT to-zone LT policy LT match destination-
address any
set logical-systems LSYS3A security policies from-zone LT to-zone LT policy LT match application
any
set logical-systems LSYS3A security policies from-zone LT to-zone LT policy LT then permit
```

```
set logical-systems LSYS3A security policies default-policy permit-all
set logical-systems LSYS3A security zones security-zone LT host-inbound-traffic system-services
all
set logical-systems LSYS3A security zones security-zone LT host-inbound-traffic protocols all
set logical-systems LSYS3A security zones security-zone LT interfaces lt-0/0/0.31
```

Step-by-Step Procedure

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

1. Define a security profile and assign to a logical system.

```
[edit]
user@host# set system security-profile SP-user logical-system LSYS3A
```

2. Set the LT interface as encapsulation frame-relay in the logical system.

```
[edit]
user@host# set logical-systems LSYS3 interfaces lt-0/0/0 unit 30 encapsulation frame-relay
```

3. Configure the logical tunnel interface by including the dlci.

```
[edit]
user@host# set logical-systems LSYS3 interfaces lt-0/0/0 unit 30 dlci 16
```

4. Configure a peer unit relationship between LT interfaces, thus creating a point-to-point connection.

```
[edit]
user@host# set logical-systems LSYS3 interfaces lt-0/0/0 unit 30 peer-unit 31
```

5. Specify the IP address for the LT interface.

```
[edit]
user@host# set logical-systems LSYS3 interfaces lt-0/0/0 unit 30 family inet address
192.255.3.1/30
```

6. Set the security zone for the LT interface.

```
[edit]
user@host# set logical-systems LSYS3 security zones security-zone LT interfaces lt-0/0/0.30
```

7. Set the LT interface as encapsulation frame-relay in the logical system.

```
[edit]
user@host# set logical-systems LSYS3A interfaces lt-0/0/0 unit 31 encapsulation frame-relay
```

8. Configure the logical tunnel interface by including the dlci.

```
[edit]
user@host# set logical-systems LSYS3A interfaces lt-0/0/0 unit 31 dlci 16
```

9. Configure a peer unit relationship between LT interfaces, thus creating a point-to-point connection.

```
[edit]
user@host# set logical-systems LSYS3A interfaces lt-0/0/0 unit 31 peer-unit 30
```

10. Specify the IP address for the LT interface.

```
[edit]
user@host# set logical-systems LSYS3A interfaces lt-0/0/0 unit 31 family inet address
192.255.3.2/30
```

11. Configure a security policy that permits traffic from the LT zone to the LT policy LT zone.

```
[edit]
user@host# set logical-systems LSYS3A security policies from-zone LT to-zone LT policy LT
match source-address any
user@host# set logical-systems LSYS3A security policies from-zone LT to-zone LT policy LT
match destination-address any
user@host# set logical-systems LSYS3A security policies from-zone LT to-zone LT policy LT
match application any
```

```
user@host# set logical-systems LSYS3A security policies from-zone LT to-zone LT policy LT
then permit
```

12. Configure a security policy that permits traffic from default-policy.

```
[edit]
user@host# set logical-systems LSYS3A security policies default-policy permit-all
```

13. Configure security zones.

```
[edit]
user@host# set logical-systems LSYS3A security zones security-zone LT host-inbound-traffic
system-services all
user@host# set logical-systems LSYS3A security zones security-zone LT host-inbound-traffic
protocols all
user@host# set logical-systems LSYS3A security zones security-zone LT interfaces
lt-0/0/0.31
```

Results

- From configuration mode, confirm your configuration by entering the `show logical-systems LSYS3` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show logical-systems LSYS3
interfaces {
    lt-0/0/0 {
        unit 30 {
            encapsulation frame-relay;
            dlci 16;
            peer-unit 31;
            family inet {
                address 192.255.3.1/30;
            }
        }
    }
}
security {
```

```
    zones {
        security-zone LT {
            interfaces {
                lt-0/0/0.30;
            }
        }
    }
}
```

- From configuration mode, confirm your configuration by entering the `show logical-systems LSYS3A` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show logical-systems LSYS3A
```

```
    interfaces {
        lt-0/0/0 {
            unit 31 {
                encapsulation frame-relay;
                dlci 16;
                peer-unit 30;
                family inet {
                    address 192.255.3.2/30;
                }
            }
        }
    }
    security {
        policies {
            from-zone LT to-zone LT {
                policy LT {
                    match {
                        source-address any;
                        destination-address any;
                        application any;
                    }
                    then {
                        permit;
                    }
```

```
                }
            }
            default-policy {
                permit-all;
            }
        }
        zones {
            security-zone LT {
                host-inbound-traffic {
                    system-services {
                        all;
                    }
                    protocols {
                        all;
                    }
                }
                interfaces {
                    lt-0/0/0.31;
                }
            }
        }
    }
```

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

Configuring Logical Systems Interconnect with Multiple VPLS Switches

CLI Quick Configuration

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

```
set interfaces lt-0/0/0 unit 11 encapsulation ethernet-vpls
set interfaces lt-0/0/0 unit 11 peer-unit 1
set interfaces lt-0/0/0 unit 12 encapsulation ethernet-vpls
set interfaces lt-0/0/0 unit 12 peer-unit 2
set interfaces lt-0/0/0 unit 13 encapsulation ethernet-vpls
set interfaces lt-0/0/0 unit 13 peer-unit 3
set interfaces lt-0/0/0 unit 23 encapsulation ethernet-vpls
set interfaces lt-0/0/0 unit 23 peer-unit 22
```

```
set interfaces lt-0/0/0 unit 25 encapsulation ethernet-vpls
set interfaces lt-0/0/0 unit 25 peer-unit 24
set routing-instances vpls-switch-1 instance-type vpls
set routing-instances vpls-switch-1 interface lt-0/0/0.11
set routing-instances vpls-switch-1 interface lt-0/0/0.12
set routing-instances vpls-switch-1 interface lt-0/0/0.13
set routing-instances vpls-switch-2 instance-type vpls
set routing-instances vpls-switch-2 interface lt-0/0/0.23
set routing-instances vpls-switch-2 interface lt-0/0/0.25
set logical-systems LSYS1 interfaces lt-0/0/0 unit 1 encapsulation ethernet
set logical-systems LSYS1 interfaces lt-0/0/0 unit 1 peer-unit 11
set logical-systems LSYS1 interfaces lt-0/0/0 unit 1 family inet address 192.255.0.1/24
set logical-systems LSYS2 interfaces lt-0/0/0 unit 2 encapsulation ethernet
set logical-systems LSYS2 interfaces lt-0/0/0 unit 2 peer-unit 12
set logical-systems LSYS2 interfaces lt-0/0/0 unit 2 family inet address 192.255.0.2/24
set logical-systems LSYS2 interfaces lt-0/0/0 unit 22 encapsulation ethernet
set logical-systems LSYS2 interfaces lt-0/0/0 unit 22 peer-unit 23
set logical-systems LSYS2 interfaces lt-0/0/0 unit 22 family inet address 192.255.4.1/30
set logical-systems LSYS3 interfaces lt-0/0/0 unit 3 encapsulation ethernet
set logical-systems LSYS3 interfaces lt-0/0/0 unit 3 peer-unit 13
set logical-systems LSYS3 interfaces lt-0/0/0 unit 3 family inet address 192.255.0.3/24
set logical-systems LSYS2B interfaces lt-0/0/0 unit 24 encapsulation ethernet
set logical-systems LSYS2B interfaces lt-0/0/0 unit 24 peer-unit 25
set logical-systems LSYS2B interfaces lt-0/0/0 unit 24 family inet address 192.255.4.2/30
set system security-profile SP-user policy maximum 100
set system security-profile SP-user policy reserved 50
set system security-profile SP-user zone maximum 60
set system security-profile SP-user zone reserved 10
set system security-profile SP-user flow-session maximum 100
set system security-profile SP-user flow-session reserved 50
set system security-profile SP-user logical-system LSYS1
set system security-profile SP-user logical-system LSYS2
set system security-profile SP-user logical-system LSYS3
set system security-profile SP-user logical-system LSYS2B
```

Step-by-Step Procedure

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

1. Configure the lt-0/0/0 interfaces.

```
[edit]
user@host# set interfaces lt-0/0/0 unit 11 encapsulation ethernet-vpls
user@host# set interfaces lt-0/0/0 unit 11 peer-unit 1
user@host# set interfaces lt-0/0/0 unit 12 encapsulation ethernet-vpls
user@host# set interfaces lt-0/0/0 unit 12 peer-unit 2
user@host# set interfaces lt-0/0/0 unit 13 encapsulation ethernet-vpls
user@host# set interfaces lt-0/0/0 unit 13 peer-unit 3
user@host# set interfaces lt-0/0/0 unit 23 encapsulation ethernet-vpls
user@host# set interfaces lt-0/0/0 unit 23 peer-unit 22
user@host# set interfaces lt-0/0/0 unit 25 encapsulation ethernet-vpls
user@host# set interfaces lt-0/0/0 unit 25 peer-unit 24
```

2. Configure the routing instance for the VPLS switches and add interfaces to it.

```
[edit]
user@host# set routing-instances vpls-switch-1 instance-type vpls
user@host# set routing-instances vpls-switch-1 interface lt-0/0/0.11
user@host# set routing-instances vpls-switch-1 interface lt-0/0/0.12
user@host# set routing-instances vpls-switch-1 interface lt-0/0/0.13
user@host# set routing-instances vpls-switch-2 instance-type vpls
user@host# set routing-instances vpls-switch-2 interface lt-0/0/0.23
user@host# set routing-instances vpls-switch-2 interface lt-0/0/0.25
```

3. Configure LSYS1 with lt-0/0/0.1 interface and peer lt-0/0/0.11.

```
[edit]
user@host# set logical-systems LSYS1 interfaces lt-0/0/0 unit 1 encapsulation ethernet
user@host# set logical-systems LSYS1 interfaces lt-0/0/0 unit 1 peer-unit 11
user@host# set logical-systems LSYS1 interfaces lt-0/0/0 unit 1 family inet address
192.255.0.1/24
```

4. Configure LSYS2 with lt-0/0/0.2 interface and peer lt-0/0/0.12.

```
[edit]
user@host# set logical-systems LSYS2 interfaces lt-0/0/0 unit 2 encapsulation ethernet
user@host# set logical-systems LSYS2 interfaces lt-0/0/0 unit 2 peer-unit 12
user@host# set logical-systems LSYS2 interfaces lt-0/0/0 unit 2 family inet address
```

```
192.255.0.2/24
user@host# set logical-systems LSYS2 interfaces lt-0/0/0 unit 22 encapsulation ethernet
user@host# set logical-systems LSYS2 interfaces lt-0/0/0 unit 22 peer-unit 23
user@host# set logical-systems LSYS2 interfaces lt-0/0/0 unit 22 family inet address
192.255.4.1/30
```

5. Configure LSYS3 with lt-0/0/0.3 interface and peer lt-0/0/0.13

```
[edit]
user@host# set logical-systems LSYS3 interfaces lt-0/0/0 unit 3 encapsulation ethernet
user@host# set logical-systems LSYS3 interfaces lt-0/0/0 unit 3 peer-unit 13
user@host# set logical-systems LSYS3 interfaces lt-0/0/0 unit 3 family inet address
192.255.0.3/24
```

6. Configure LSYS2B with lt-0/0/0 interface and peer-unit 24.

```
[edit]
user@host# set logical-systems LSYS2B interfaces lt-0/0/0 unit 24 encapsulation ethernet
user@host# set logical-systems LSYS2B interfaces lt-0/0/0 unit 24 peer-unit 25
user@host# set logical-systems LSYS2B interfaces lt-0/0/0 unit 24 family inet address
192.255.4.2/30
```

7. Assign security-profile for logical-systems.

```
[edit]
user@host# set system security-profile SP-user policy maximum 100
user@host# set system security-profile SP-user policy reserved 50
user@host# set system security-profile SP-user zone maximum 60
user@host# set system security-profile SP-user zone reserved 10
user@host# set system security-profile SP-user flow-session maximum 100
user@host#set system security-profile SP-user flow-session reserved 50
user@host# set system security-profile SP-user logical-system LSYS1
user@host# set system security-profile SP-user logical-system LSYS2
user@host# set system security-profile SP-user logical-system LSYS3
user@host# set system security-profile SP-user logical-system LSYS2B
```

Results

- From configuration mode, confirm your configuration by entering the `show interfaces lt-0/0/0`, command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it

```
[edit]
user@host# show interfaces lt-0/0/0
unit 11 {
    encapsulation ethernet-vpls;
    peer-unit 1;
}
unit 12 {
    encapsulation ethernet-vpls;
    peer-unit 2;
}
unit 13 {
    encapsulation ethernet-vpls;
    peer-unit 3;
}
unit 23 {
    encapsulation ethernet-vpls;
    peer-unit 22;
}
unit 25 {
    encapsulation ethernet-vpls;
    peer-unit 24;
}
```

- From configuration mode, confirm your configuration by entering the `show routing-instances`, command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show routing-instances
vpls-switch-1 {
    instance-type vpls;
    interface lt-0/0/0.11;
    interface lt-0/0/0.12;
    interface lt-0/0/0.13;
}
```

```
vpls-switch-2 {
    instance-type vpls;
    interface lt-0/0/0.23;
    interface lt-0/0/0.25;
}
```

- From configuration mode, confirm your configuration by entering the `show logical-systems LSYS1`, command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show logical-systems LSYS1
interfaces {
    lt-0/0/0 {
        unit 1 {
            encapsulation ethernet;
            peer-unit 11;
            family inet {
                address 192.255.0.1/24;
            }
        }
    }
}
```

- From configuration mode, confirm your configuration by entering the `show logical-systems LSYS2`, command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show logical-systems LSYS2
interfaces {
    lt-0/0/0 {
        unit 2 {
            encapsulation ethernet;
            peer-unit 12;
            family inet {
                address 192.255.0.2/24;
            }
        }
        unit 22 {
            encapsulation ethernet;
```

```
            peer-unit 23;
            family inet {
                address 192.255.4.1/30;
            }
        }
    }
}
```

- From configuration mode, confirm your configuration by entering the `show logical-systems LSYS3`, command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show logical-systems LSYS3
interfaces {
    lt-0/0/0 {
        unit 3 {
            encapsulation ethernet;
            peer-unit 13;
            family inet {
                address 192.255.0.3/24;
            }
        }
    }
}
```

- From configuration mode, confirm your configuration by entering the `show logical-systems LSYS2B`, command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show logical-systems LSYS2B
interfaces {
    lt-0/0/0 {
        unit 24 {
            encapsulation ethernet;
            peer-unit 25;
            family inet {
                address 192.255.4.2/30;
            }
        }
```

```
    }
}
```

- From configuration mode, confirm your configuration by entering the `show system security-profile`, command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show system security-profile
SP-user {
    policy {
        maximum 100;
        reserved 50;
    }
    zone {
        maximum 60;
        reserved 10;
    }
    flow-session {
        maximum 100;
        reserved 50;
    }
    logical-system [ LSYS1 LSYS2 LSYS3 LSYS2B ];
}
```

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

Verification

IN THIS SECTION

- Verifying the Security-Profile for all Logical-systems | 49
- Verifying the LT Interfaces for all Logical systems | 49

To confirm that the configuration is working properly, perform these tasks:

Verifying the Security-Profile for all Logical-systems

Purpose

Verify security profile for each logical systems.

Action

From operational mode, enter the `show system security-profile security-log-stream-number logical-system all` command.

```
user@host> show system security-profile security-log-stream-number logical-system all
```

```
logical system name   security profile name       usage    reserved   maximum
root-logical-system   Default-Profile                2           0          2000
LSYS1                 SP-user                        1          10          60
LSYS2                 SP-user                        1          10          60
LSYS2B                SP-user                        1          10          60
LSYS3                 SP-user                        1          10          60
```

Meaning

The output provides the usage and reserved values for the logical systems when security-log-stream is configured.

Verifying the LT Interfaces for all Logical systems

Purpose

Verify interfaces for logical systems.

Action

From operational mode, enter the `show interfaces lt-0/0/0 terse` command.

```
user@host> show interfaces lt-0/0/0 terse
```

```
Interface               Admin Link Proto    Local                 Remote
lt-0/0/0                up    up
lt-0/0/0.1              up    up   inet     192.255.0.1/24
lt-0/0/0.2              up    up   inet     192.255.0.2/24
lt-0/0/0.3              up    up   inet     192.255.0.3/24
lt-0/0/0.11             up    up   vpls
lt-0/0/0.12             up    up   vpls
lt-0/0/0.13             up    up   vpls
lt-0/0/0.22             up    up   inet     192.255.4.1/30
lt-0/0/0.23             up    up   vpls
lt-0/0/0.24             up    up   inet     192.255.4.2/30
lt-0/0/0.25             up    up   vpls
lt-0/0/0.32767          up    up
```

Meaning

The output provides the status of LT interfaces. All the LT interfaces are up.

SEE ALSO

- Understanding the Primary Logical Systems and the Primary Administrator Role | 23
- Understanding User Logical Systems and the User Logical System Administrator Role | 53
- Understanding the Interconnect Logical System and Logical Tunnel Interfaces | 11
- SRX Series Logical Systems Primary Administrator Configuration Tasks Overview | 24

User Logical Systems Overview

IN THIS SECTION

- User Logical Systems Configuration Overview | 51
- Understanding User Logical Systems and the User Logical System Administrator Role | 53

A user logical system enables you to configure zones, security policies, logical interfaces and security resources assigned to its own user logical system. For more information, see the following topics:

User Logical Systems Configuration Overview

When the primary administrator creates a user logical system, he assigns a user logical system administrator to manage it. A user logical system can have multiple user logical system administrators.

As a user logical system administrator, you can access and view resources in your user logical system but not those of other user logical systems or the primary logical system. You can configure resources allocated to your user logical system, but you cannot modify the numbers of allocated resources.

The following procedure lists the tasks that the user logical system administrator performs to configure resources in the user logical system:

1. Log in to the user logical system with the login and password configured by the primary administrator:

   a. SSH to the management IP address configured on the device. Log in to the user logical system with the administrator login and password provided by the primary administrator.

      You enter a UNIX shell in the user logical system configured by the primary administrator.

      Starting in Junos OS Release 20.1R1, On SRX5400, SRX5600, and SRX5800 Series devices, Trusted Platform Module (TPM) supports only with SRX5K-RE3-128G Routing Engine (RE3). The TPM chip enables by default. To use the TPM functionality in logical systems, you must configure Master-Encryption-key (MEK) at the root logical system only, and the user logical systems will inherit the same MEK to encrypt configuration hash and public key infrastructure (PKI) key-pairs. For more information on TPM, see *Using Trusted Platform Module to Bind Secrets on SRX Series Devices*.

b. The presence of the > prompt indicates the CLI has started. The prompt is preceded by a string that contains your username, the hostname of the router, and the name of the user logical system. When the CLI starts, you are at the top level in operational mode. You enter configuration mode by entering the **configure** operational mode command. The CLI prompt changes from user@host: *logical-system*> to user@host: *logical-system*#.

   To exit the CLI and return to the UNIX shell, enter the **quit** command.

2. Configure the logical interfaces assigned to the user logical system by the primary administrator. Configure one or more routing instances and the routing protocols and options within each instance. See "Example: Configuring Interfaces and Routing Instances for a User Logical Systems" on page 139.

3. Configure security resources for the user logical system:

   a. Create zones for the user logical system and bind the logical interfaces to the zones. Address books can be created that are attached to zones for use in policies. See "Example: Configuring Security Zones for a User Logical Systems" on page 167.

   b. Configure screen options at the zone level. See "Example: Configuring Screen Options for a User Logical Systems" on page 224.

   c. Configure security policies between zones in the user logical system. See "Example: Configuring Security Policies in a User Logical Systems" on page 216.

      Custom applications or application sets can be created for specific types of traffic. To create a custom application, use the `application` configuration statement at the [`edit applications`] hierarchy level. To create an application set, use the `application-set` configuration statement at the [`edit applications`] hierarchy level.

   d. Configure firewall authentication. The primary administrator creates access profiles in the primary logical system. See "Example: Configuring Access Profiles (Primary Administrators Only)" on page 173.

      The user logical system administrator then configures a security policy that specifies firewall authentication for matching traffic and configures the type of authentication (pass-through or Web authentication), default access profile, and success banner. See "Example: Configuring Firewall Authentication for a User Logical System" on page 187.

   e. Configure a route-based VPN tunnel to secure traffic between a user logical system and a remote site. The primary administrator assigns a secure tunnel interface to the user logical system and configures IKE and IPsec SAs for the VPN tunnel. See "Example: Configuring IKE and IPsec SAs for a VPN Tunnel (Primary Administrators Only)" on page 235.

      The user logical system administrator then configures a route-based VPN tunnel. See "Example: Configuring a Route-Based VPN Tunnel in a User Logical Systems" on page 245.

f. Configure Network Address Translation (NAT). See "Example: Configuring Network Address Translation for a User Logical Systems" on page 134.

g. Configure and assigning a predefined IDP policy to the user logical system. The primary administrator configures IDP policies at the root level and specifies an IDP policy in the security profile that is bound to a logical system. See "Example: Configuring and Assigning a Predefined IDP Policy for a User Logical System" on page 288.

   The user logical system administrator then enables IDP in a security policy. See "Example: Enabling IDP in a User Logical System Security Policy" on page 292.

h. Configure and enable an IDP policy at the user logical system. See "Example: Configuring an IDP Policy for a User Logical System" on page 296

i. Display or clear application system cache (ASC) entries. See "Understanding Logical Systems Application Identification Services" on page 339.

j. Configure application firewall services on a user logical system. See "Understanding Logical Systems Application Firewall Services" on page 341 and "Example: Configuring Application Firewall Services for a User Logical System" on page 349.

k. Configure the AppTrack application tracking tool. See "Example: Configuring AppTrack for a User Logical Systems" on page 355.

SEE ALSO

Example: Configuring User Logical Systems | 150

Understanding User Logical Systems and the User Logical System Administrator Role | 53

Understanding User Logical Systems and the User Logical System Administrator Role

Logical systems allow a primary administrator to partition an SRX Series device into discrete contexts called user logical systems. User logical systems are self-contained, private contexts, separate both from one another and from the primary logical system. A user logical system has its own security, networking, logical interfaces, routing configurations, and one or more user logical system administrators.

When the primary administrator creates a user logical system, he assigns one or more user logical system administrators to manage it. A user logical system administrator has a view of the device that is limited to his logical system. Although a user logical system is managed by a user logical system

administrator, the primary administrator has a global view of the device and access to all user logical systems. If necessary, the primary administrator can manage any user logical system on the device.

The role and responsibilities of a user logical system administrator differ from those of the primary administrator. As a user logical system administrator, you can access, configure, and view the configuration for your user logical system resources, but not those of other user logical systems or the primary logical system.

As a user logical system administrator, you can:

- Configure zones, address books, security policies, user lists, custom services, and so forth, for your user logical system environment, based on the resources allocated to it.

  For example, if the primary administrator allocates 40 zones to your user logical system, you can configure and administer those zones, but you cannot change the allocated number.

- Configure routing instances and assign allotted interfaces to them. Create static routes and add them to your routing instances. Configure routing protocols.

- Configure, enable, and monitor application firewall policy on your user logical system.

- Configure AppTrack.

- View all assigned logical interfaces and configure their attributes. The attributes that you configure for logical interfaces for your user logical system cannot be seen by other user logical system administrators.

- Run operational commands for your user logical system.

SEE ALSO

Understanding Logical Systems for SRX Series Services Gateways | **5**

Example: Configuring Interfaces, Routing Instances, and Static Routes for the Primary and Interconnect Logical Systems and Logical Tunnel Interfaces for the User Logical Systems (Primary Administrators Only) | **114**

Understanding Logical Systems Security Profiles (Primary Administrators Only) | **71**

Example: Configuring Logical Systems Security Profiles (Primary Administrators Only) | **77**

Setting Up a Logical System

IN THIS SECTION

- Example: Configuring Root Password for Logical Systems | 55
- Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System | 57

Example: Configuring Root Password for Logical Systems

IN THIS SECTION

- Requirements | 55
- Overview | 55
- Configuration | 56

Requirements

Before you begin, read "SRX Series Logical Systems Primary Administrator Configuration Tasks Overview" on page 24 to understand how this task fits into the overall configuration process.

The example uses an SRX5600 device running Junos OS with logical systems.

Overview

IN THIS SECTION

- Topology | 56

The Junos OS software is installed on the router before it is delivered from the factory. When you power on your router, it is ready for you to configure. Initially you log in as *root* user without using a password.

After you log in, you can configure a password for the root user, or, in logical systems terms, the primary administrator. The primary administrator has root privileges over the device.

Topology

Configuration

IN THIS SECTION

- Configuring the Root Password | 56

Configuring the Root Password

Step-by-Step Procedure

- Configure a root password for the device.

```
user@host# set system root-authentication Talk22rt6
```

SEE ALSO

Understanding the Primary Logical Systems and the Primary Administrator Role | 23

Understanding Logical Systems for SRX Series Services Gateways | 5

Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System

IN THIS SECTION

- Requirements | 57
- Overview | 57
- Configuration | 59
- Verification | 67

This example shows how to create user logical systems and assign administrators to them. It shows how to add users to a user logical system. And the example shows how to create an interconnect logical system, which is optional.

NOTE: Only the primary administrator can create user login accounts for administrators and users. If a user logical system administrator wants to add users to his logical system, he must convey the information to the primary administrator, who will add the users.

Requirements

The example uses an SRX5600 device running Junos OS with logical systems.

Overview

IN THIS SECTION

- Topology | 58

Before you begin, read "SRX Series Logical Systems Primary Administrator Configuration Tasks Overview" on page 24 to understand how this task fits into the overall configuration process.

This example is for a company that includes product design, marketing, and accounting departments. The company wants to curtail hardware and energy costs, but not at the risk of exposing data across departments or to the Internet.

Each department has its own security requirements in regard both to other departments and to the Internet. To meet its requirements for cost control without forfeiting security, the company deploys the SRX5600 device. The primary administrator configures three user logical systems giving each department a logical device that is private and fully secured.

This topic covers how to:

- Create user logical systems and an interconnect logical system that is used as an internal VPLS switch to allow traffic to pass from one logical system to another.
- Create administrators for user logical systems other than the interconnect logical system. A user logical system can have more than one administrator. The interconnect logical system does not require an administrator.
- Add users to a user logical system.

> **NOTE**: This example shows how to configure only two users—lsdesignuser1 and lsdesignuser2. In reality, every user logical system will include many users that would require configurations similar to those shown in this example.

Topology

Figure 5 on page 59 shows an SRX5600 device deployed and configured for logical systems. The configuration examples reflect this deployment.

Figure 5: SRX Series Device Configured for Logical Systems

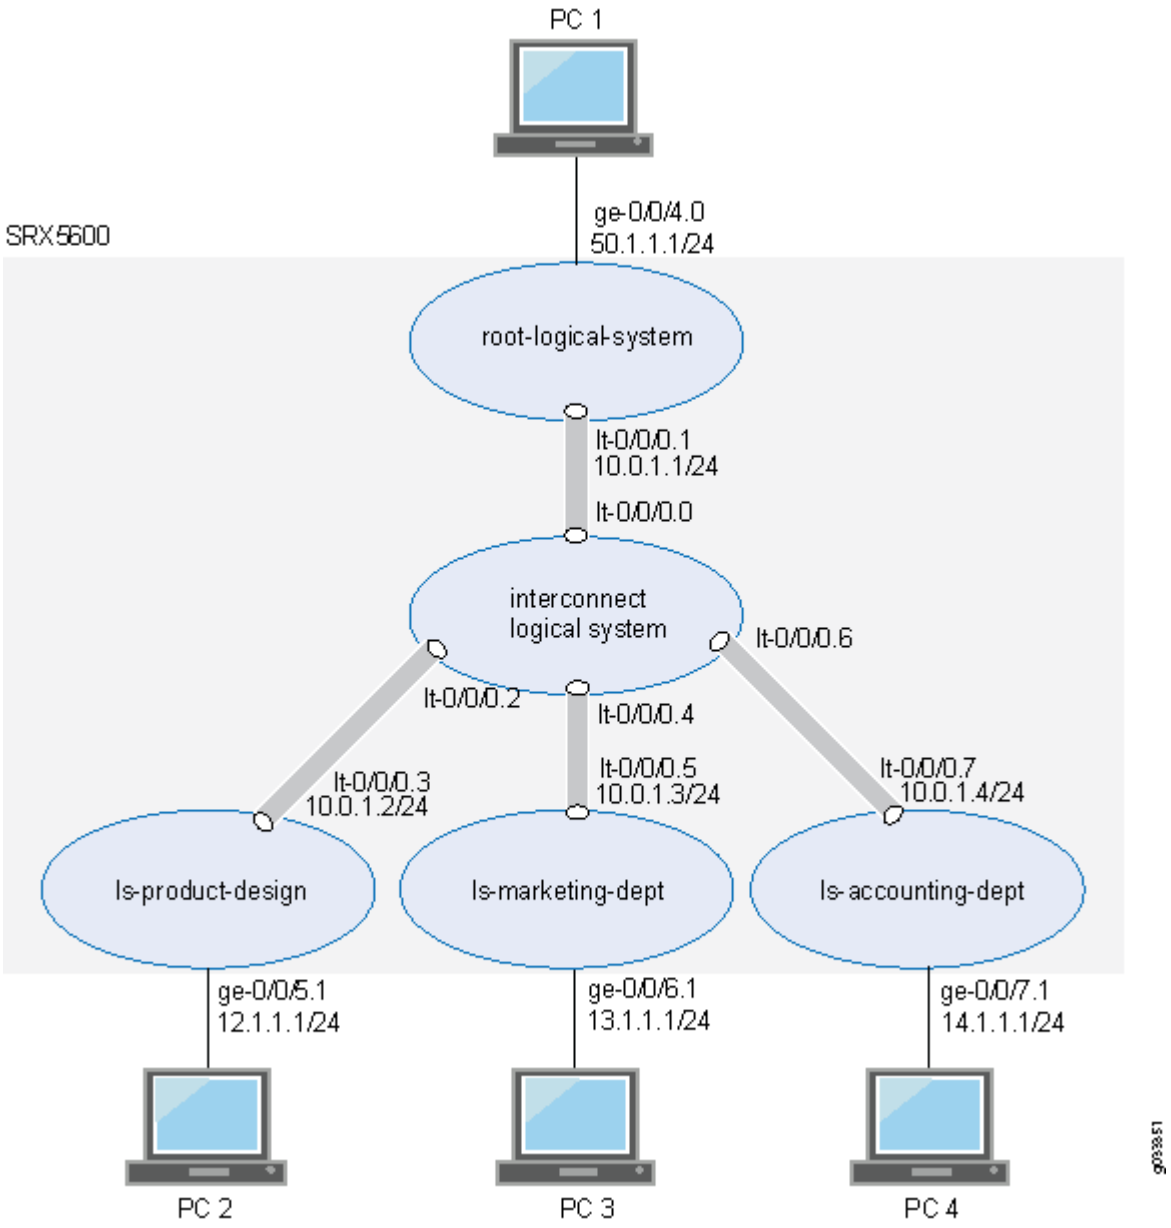

Configuration

IN THIS SECTION

- Configuring User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System | **60**

Configuring User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System

CLI Quick Configuration

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

```
set logical-systems ls-product-design
set system login class ls-design-admin logical-system ls-product-design
set system login class ls-design-admin permissions all
set system login user lsdesignadmin1 full-name lsdesignadmin1
set system login user lsdesignadmin1 class ls-design-admin
set system login user lsdesignadmin1 authentication encrypted-password "$ABC123"
set system login class ls-design-user logical-system ls-product-design
set system login class ls-design-user permissions view
set system login user lsdesignuser1 full-name lsdesignuser1
set system login user lsdesignuser1 class ls-design-user
set system login user lsdesignuser1 authentication encrypted-password "$ABC123"
set system login user lsdesignuser2 full-name lsdesignuser2
set system login user lsdesignuser2 class ls-design-user
set system login user lsdesignuser2 authentication encrypted-password "$ABC123"
set logical-systems ls-marketing-dept
set system login class ls-marketing-admin logical-system ls-marketing-dept
set system login class ls-marketing-admin permissions all
set system login user lsmarketingadmin1 class ls-marketing-admin
set system login user lsmarketingadmin1 full-name lsmarketingadmin1
set system login user lsmarketingadmin1 authentication encrypted-password "$ABC123"
set system login user lsmarketingadmin2 full-name lsmarketingadmin2
set system login user lsmarketingadmin2 class ls-marketing-admin
set system login user lsmarketingadmin2 authentication encrypted-password "$ABC123"
set logical-systems ls-accounting-dept
set system login class ls-accounting-admin logical-system ls-accounting-dept
set system login class ls-accounting-admin permissions all
set system login user lsaccountingadmin1 full-name lsaccountingadmin1
set system login user lsaccountingadmin1 class ls-accounting-admin
set system login user lsaccountingadmin1 authentication encrypted-password "$ABC123"
set logical-systems interconnect-logical-system
```

Step-by-Step Procedure

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

1. Create the first user logical system and define its administrator.

   #### Step-by-Step Procedure

   a. Create the user logical system.

   ```
   [edit]
   user@host# set logical-systems ls-product-design
   ```

   b. Assign the user login class to the user logical system.

   ```
   [edit system]
   user@host# set login class ls-design-admin logical-system ls-product-design
   ```

   c. Create the login class to give the user logical system administrator full permission over the user logical system.

   ```
   [edit system]
   user@host# set login class ls-design-admin permissions all
   ```

   d. Assign a full name to the user logical system administrator.

   ```
   [edit system]
   user@host# set login user lsdesignadmin1 full-name lsdesignadmin1
   ```

   e. Associate the login class with the user logical system administrator to allow the administrator to log in to the user logical system.

   ```
   [edit system]
   user@host# set login user lsdesignadmin1 class ls-design-admin
   ```

f. Create a user login password for the user logical system administrator.

```
[edit system]
user@host# set login user lsdesignadmin1 authentication plain-text-password
New password: Talk1234
Retype new password: Talk1234
```

2. Configure the first user for the logical system.

Step-by-Step Procedure

a. Configure the user login class and assign it to the user logical system.

```
[edit system]
user@host# set login class ls-design-user logical-system ls-product-design
```

b. To give the first user the ability to see the logical system's resources and settings but not change them, assign `view` as the permission to the login class.

```
[edit system]
user@host# set login class ls-design-user permissions view
```

c. Assign a full name to the logical system user.

```
[edit system]
user@host# set login user lsdesignuser1 full-name lsdesignuser1
```

d. Associate the login class with the user to allow the user to log in to the user logical system.

```
user@host# set login user lsdesignuser1 class ls-design-user
```

e. Create a user login password for the user.

```
[edit system]
user@host# set login user lsdesignuser1 authentication plain-text-password
```

```
New password: Talk4234
Retype new password: Talk4234
```

3. Create the second user for logical system ls-product-design.

   Step-by-Step Procedure

   a. Assign a full name to the user.

   ```
   [edit system]
   user@host# set login user lsdesignuser2 full-name lsdesignuser2
   ```

   b. Associate the user with the login class to allow the user to log in to the user logical system.

   ```
   user@host# set login user lsdesignuser2 class ls-design-user
   ```

   c. Create a user login password.

   ```
   [edit system]
   user@host# set login user lsdesignuser2 authentication plain-text-password
   New password: Talk9234
   Retype new password: Talk9234
   ```

4. Create the second user logical system and define its administrator.

   Step-by-Step Procedure

   a. Create the user logical system.

   ```
   [edit]
   user@host# set logical-systems ls-marketing-dept
   ```

   b. Configure the user login class and assign it to the user logical system.

   ```
   [edit system]
   user@host# set login class ls-marketing-admin logical-system ls-marketing-dept
   ```

c. To give the user logical system administrator control over the user logical system, assign `all` as the permissions to the login class.

```
[edit system]
user@host# set login class ls-marketing-admin permissions all
```

d. Assign a full name to the user logical system administrator.

```
[edit system]
user@host# set login user lsmarketingadmin1 full-name lsmarketingadmin1
```

e. Associate the user logical system administrator with the login class to allow the administrator to log in to the user logical system.

```
[edit system]
user@host# set login user lsmarketingadmin1 class ls-marketing-admin
```

f. Create a user login password for the user logical system administrator.

```
[edit system]
user@host# set login user lsmarketingadmin1 authentication plain-text-password
New password: Talk2345
Retype new password: Talk2345
```

5. Create a second user logical system administrator for the ls-marketing-dept logical system.

Step-by-Step Procedure

a. Assign a full name to the user logical system administrator.

```
[edit system]
user@host# set login user lsmarketingadmin2 full-name lsmarketingadmin2
```

b. Associate the user logical system administrator with the login class to allow the administrator to log in to the user logical system.

```
[edit system]
user@host# set login lsmarketingadmin2 class ls-marketing-admin
```

c. Create a user login password for the user logical system administrator.

```
[edit system]
user@host# set login user lsmarketingadmin2 authentication plain-text-password
New password: Talk6345
Retype new password: Talk6345
```

6. Create the third user logical system and define its administrator.

Step-by-Step Procedure

a. Create the user logical system.

```
[edit]
user@host# set logical-systems ls-accounting-dept
```

b. Configure the user login class and assign it to the user logical system.

```
[edit system]
user@host# set login class ls-accounting-admin logical-system ls-accounting-dept
```

c. To give the user logical system administrator control over the user logical system, assign permissions to the login class.

```
[edit system]
user@host# set login class ls-accounting-admin permissions all
```

d. Assign a full name to the user logical system administrator.

```
[edit system]
user@host# set login user lsaccountingadmin1 full-name lsaccountingadmin1
```

e. Associate the user logical system administrator with the login class to allow the administrator to log in to the user logical system.

```
[edit system]
user@host# set login user lsaccountingadmin1 class ls-accounting-admin
```

f. Create a login password for the user logical system administrator.

```
[edit system]
user@host# set login user lsaccountingadmin1 authentication plain-text-password
New password: Talk5678
Retype new password: Talk5678
```

7. Configure an interconnect logical system to allow logical systems to pass traffic from one to another.

```
user@host# set logical-systems interconnect-logical-system
```

Results

From configuration mode, confirm your configuration by entering the `show logical-systems` command to verify that the logical systems were created. Also enter the `show system login class` command for each class that you defined.

To ensure that the logical systems administrators were created, enter the `show system login user` command.

If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show logical-systems ?
interconnect-logical-system;
ls-accounting-dept;
```

```
ls-marketing-dept;
ls-product-design;
```

```
user@host# show system login class ls-design-admin
logical-system ls-product-design;
permissions all;
```

```
user@host# show system login class ls-design-user
logical-system ls-product-design
permissions view;
```

```
user@host show system login class ls-marketing-admin
logical-system ls-marketing-dept;
permissions all;
```

```
user@host show system login class ls-accounting-admin
 logical-system ls-accounting-dept;
permissions all;
```

```
user@host show system login user ?
 lsaccountingadmin1   lsaccountingadmin1
lsdesignadmin1       lsdesignadmin1
lsdesignuser2        lsdesignuser2
lsmarketingadmin1    lsmarketingadmin1
lsmarketingadmin2    lsmarketingadmin2
```

Verification

IN THIS SECTION

- Verifying User Logical Systems and Login Configurations from the Primary Logical System | 68
- Verifying User Logical Systems and Login Configurations Using SSH | 69

To confirm that the configuration is working properly, perform these tasks:

Verifying User Logical Systems and Login Configurations from the Primary Logical System

Purpose

Verify that the user logical systems exist and that you, as the primary administrator, can enter them from root. Return from a user logical system to the primary logical system.

Action

From operational mode, enter the following command:

```
root@host> set cli logical-system ls-product-design
Logical system:ls-product-design
root@host:ls-product-design>
```

```
root@host:ls-product-design> clear cli logical-system
Cleared default logical system
root@host>
```

```
root@host> set cli logical-system ls-marketing-dept
Logical system:ls-marketing-dept
root@host:ls-marketing-dept>
```

```
root@host:ls-marketing-dept> clear cli logical-system
Cleared default logical system
root@host>
```

```
root@host> set cli logical-system ls-accounting-dept
Logical system:ls-accounting-dept
root@host:ls-accounting-dept>
```

```
root@host:ls-accounting-dept> clear cli logical-system
Cleared default logical system
root@host>
```

Verifying User Logical Systems and Login Configurations Using SSH

Purpose

Verify that the user logical systems you created exist and that the administrators' login IDs and passwords that you created are correct.

Action

Use SSH to log in to each user logical system as its user administrator would do.

1. Run SSH specifying the IP address of your SRX Series device.
2. Enter the login ID and password for the administrator for one of the user logical systems that you created. After you log in, the prompt shows the administrator name. Notice how this result differs from the result produced when you log in to the user logical system from the primary logical system at root. Repeat this procedure for all of your user logical systems.

```
login: lsdesignadmin1
Password: Talk1234

lsdesignadmin1@host: ls-product-design>
```

SEE ALSO

Example: Configuring Logical Systems Security Profiles (Primary Administrators Only) | 77

Example: Configuring Interfaces, Routing Instances, and Static Routes for the Primary and Interconnect Logical Systems and Logical Tunnel Interfaces for the User Logical Systems (Primary Administrators Only) | 114

RELATED DOCUMENTATION

Understanding the Primary Logical Systems and the Primary Administrator Role | 23

Understanding Logical Systems for SRX Series Services Gateways | 5

Security Profiles for Logical Systems

IN THIS SECTION

- Understanding Logical Systems Security Profiles (Primary Administrators Only) | 71
- Example: Configuring Logical Systems Security Profiles (Primary Administrators Only) | 77

- Example: Configuring User Logical Systems Security Profiles | 89
- Example: Configuring Security log stream for Logical Systems | 96

Security profiles for logical systems allow you to allocate resources. Security profiles specifies the number of resources to allocate to a logical system to which the security profile is bound. All system resources are allocated to primary logical system and the primary administrator allocates them to user logical system using security profile. For more information, see the following topics:

Understanding Logical Systems Security Profiles (Primary Administrators Only)

IN THIS SECTION

- Logical Systems Security Profiles | 72
- How the System Assesses Resources Assignment and Use Across Logical Systems | 72
- Cases: Assessments of Reserved Resources Assigned Through Security Profiles | 75

Logical systems allow you to virtually divide a supported SRX Series device into multiple devices, isolating one from another, securing them from intrusion and attacks, and protecting them from faulty conditions outside their own contexts. To protect logical systems, security resources are configured in a manner similar to how they are configured for a discrete device. However, as the primary administrator, you must allocate the kinds and amounts of security resources to logical systems. The logical system administrator allocates resources for his own logical system.

An SRX Series device running logical systems can be partitioned into user logical systems, an interconnect logical system, if desired, and the default primary logical system. When the system is initialized, the primary logical system is created at the root level. All system resources are assigned to it, effectively creating a default primary logical system security profile. To distribute security resources across logical systems, the primary administrator creates security profiles that specify the kinds and amounts of resources to be allocated to a logical system that the security profile is bound to. Only the primary administrator can configure security profiles and bind them to logical systems. The user logical system administrator configures these resources for his or her logical system.

Logical systems are defined largely by the resources allocated to them, including security components, interfaces, routing instances, static routes, and dynamic routing protocols. When the primary administrator configures a user logical system, he binds a security profile to it. Any attempt to commit a configuration for a user logical system without a security profile bound to it will fail.

This topic includes the following sections:

Logical Systems Security Profiles

As primary administrator, you can configure a single security profile to assign resources to a specific logical system, use the same security profile for more than one logical system, or use a mix of both methods. You can configure up to 32 security profiles on an SRX Series device running logical systems. When you reach the limit, you must delete a security profile and commit the configuration change before you can create and commit another security profile. In many cases fewer security profiles are needed because you might bind a single security profile to more than one logical system.

Security profiles allow you to:

- Share the device's resources, including policies, zones, addresses and address books, flow sessions, and various forms of NAT, among all logical systems appropriately. You can dedicate various amounts of a resource to the logical systems and allow them to compete for use of the free resources.

  Security profiles protect against one logical system exhausting a resource that is required at the same time by other logical systems. Security profiles protect critical system resources and maintain a fair level of performance among user logical systems when the device is experiencing heavy traffic flow. They defend against one user logical system dominating the use of resources and depriving other user logical systems of them.

- Configure the device in a scalable way to allow for future creation of additional user logical systems.

You must delete a logical system's security profile before you delete that logical system.

How the System Assesses Resources Assignment and Use Across Logical Systems

To provision a logical system with security resources, you, as a primary administrator, configure a security profile that specifies for each resource:

- A reserved quota that guarantees that the specified resource amount is always available to the logical system.

- A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota

does not guarantee that the amount specified for the resource in the security profile is available. Logical systems must compete for global resources.

If a reserved quota is not configured for a resource, the default value is 0. If a maximum allowed quota is not configured for a resource, the default value is the global system quota for the resource (global system quotas are platform-dependent). The primary administrator must configure appropriate maximum allowed quota values in the security profiles so the maximum resource usage of a specific logical system does not negatively impact other logical systems configured on the device. The primary administrator must configure the appropriate maximum-allowed quota values in the security profiles so that the maximum resource usage of a specific logical system does not negatively impact other logical systems configured on the device.

The system maintains a count of all allocated resources that are reserved, used, and made available again when a logical system is deleted. This count determines whether resources are available to use for new logical systems or to increase the amount of the resources allocated to existing logical systems through their security profiles.

When a user logical system is deleted, its reserved resource allocations are released for use by other logical systems.

Resources configured in security profiles are characterized as static modular resources or dynamic resources. For static resources, we recommend setting a maximum quota for a resource equal or close to the amount specified as its reserved quota, to allow for scalable configuration of logical systems. A high maximum quota for a resource might give a logical system greater flexibility through access to a larger amount of that resource, but it would constrain the amount available to allocate to a new user logical system.

The difference between reserved and maximum allowed amounts for a dynamic resource is not important because dynamic resources are aged out and do not deplete the pool available for assignment to other logical systems.

The following resources can be specified in a security profile:

- Security policies, including schedulers
- Security zones
- Addresses and address books for security policies
- Application firewall rule sets
- Application firewall rules
- Firewall authentication
- Flow sessions and gates
- NAT, including:

- Cone NAT bindings
- NAT destination rule
- NAT destination pool
- NAT IP address in source pool without Port Address Translation (PAT)

  NOTE: IPv6 addresses in IPv6 source pools without PAT are not included in security profiles.

- NAT IP address in source pool with PAT
- NAT port overloading
- NAT source pool
- NAT source rule
- NAT static rule

NOTE: All resources except flow sessions are static.

You can modify a logical system security profile dynamically while the security profile is assigned to other logical systems. However, to ensure that the system resource quota is not exceeded, the system takes the following actions:

- If a static quota is changed, system daemons that maintain logical system counts for resources specified in security profiles revalidate the security profile. This check identifies the number of resources assigned across all logical systems to determine whether the allocated resources, including their increased amounts, are available.

  These quota checks are the same quota checks that the system performs when you add a new user logical system and bind a security profile to it. They are also performed when you bind a different security profile from the security profile that is presently assigned to it to an existing user logical system (or the primary logical system).

- If a dynamic quota is changed, no check is performed, but the new quota is imposed on future resource usage.

Cases: Assessments of Reserved Resources Assigned Through Security Profiles

To understand how the system assesses allocation of reserved resources through security profiles, consider the following three cases that address allocation of one resource, zones. To keep the example simple, 10 zones are allocated in security-profile-1: 4 reserved zones and 6 maximum zones. This example assumes that the full maximum amount specified–six zones–is available for the user logical systems. The system maximum number of zones is 10.

These cases address configuration across logical systems. They test to see whether a configuration will succeed or fail when it is committed based on allocation of zones.

Table 2 on page 75 shows the security profiles and their zone allocations.

Table 2: Security Profiles Used for Reserved Resource Assessments

Two Security Profiles Used in the Configuration Cases
security-profile-1 • zones reserved quota = 4 • zones maximum quota = 6 **NOTE**: Later the primary administrator dynamically increases the reserved zone count specified in this profile.
primary-logical-system-profile • zones maximum quota = 10 • no reserved quota

Table 3 on page 76 shows three cases that illustrate how the system assesses reserved resources for zones across logical systems based on security profile configurations.

- The configuration for the first case succeeds because the cumulative reserved resource quota for zones configured in the security profiles bound to all logical systems is 8, which is less than the system maximum resource quota.
- The configuration for the second case fails because the cumulative reserved resource quota for zones configured in the security profiles bound to all logical systems is 12, which is greater than the system maximum resource quota.
- The configuration for the third case fails because the cumulative reserved resource quota for zones configured in the security profiles bound to all logical systems is 12, which is greater than the system maximum resource quota.

Table 3: Reserved Resource Allocation Assessment Across Logical Systems

Reserved Resource Quota Checks Across Logical Systems
Example 1: Succeeds This configuration is within bounds: 4+4+0=8, maximum capacity =10. Security Profiles Used • The security profile security-profile-1 is bound to two user logical systems: user-logical-system-1 and user-logical-system-2. • The primary-logical-system-profile profile is used exclusively for the primary logical system. • user-logical-system-1 = 4 reserved zones. • user-logical-system-2 = 4 reserved zones. • primary-logical-system = 0 reserved zones.
Example 2: Fails This configuration is out of bounds: 4+4+4=12, maximum capacity =10. • user-logical-system-1 = 4 reserved zones. • user-logical-system-2 = 4 reserved zones. • primary-logical-system = 0 reserved zones. • new-user-logical-system = 4 reserved zones. Security Profiles • The security profile security-profile-1 is bound to two user logical systems: user-logical-system-1 and user-logical-system-2. • The primary-logical-system-profile is bound to the primary logical system and used exclusively for it. • The primary administrator configures a new user logical system called new-user-logical-system and binds security-profile-1 to it.

Table 3: Reserved Resource Allocation Assessment Across Logical Systems *(Continued)*

Reserved Resource Quota Checks Across Logical Systems
Example 3: Fails This configuration is out of bounds: 6+6=12, maximum capacity =10. The primary administrator modifies the reserved zones quota in security-profile-1, increasing the count to 6. • user-logical-system-1 = 6 reserved zones. • user-logical-system-2 = 6 reserved zones. • primary-logical-system = 0 reserved zones.

SEE ALSO

Example: Configuring Logical Systems Security Profiles (Primary Administrators Only) | 77
Understanding the Primary Logical Systems and the Primary Administrator Role | 23
Understanding User Logical Systems and the User Logical System Administrator Role | 53

Example: Configuring Logical Systems Security Profiles (Primary Administrators Only)

IN THIS SECTION

Requirements | 78
Overview | 78
Configuration | 78
Verification | 88

This example shows how a primary administrator configures three logical system security profiles to assign to user logical systems and the primary logical system to provision them with security resources.

Requirements

The example uses an SRX5600 device running Junos OS with logical systems.

Before you begin, read "SRX Series Logical Systems Primary Administrator Configuration Tasks Overview" on page 24 to understand how this task fits into the overall configuration process.

Overview

IN THIS SECTION

- Topology | 78

This example shows how to configure security profiles for the following logical systems:

- The root-logical-system logical system. The security profile primary-profile is assigned to the primary, or root, logical system.
- The ls-product-design logical system. The security profile ls-design-profile is assigned to the logical system.
- The ls-marketing-dept logical system. The security profile ls-accnt-mrkt-profile is assigned to the logical system.
- The ls-accounting-dept logical system. The security profile ls-accnt-mrkt-profile is assigned to the logical system.
- The interconnect-logical-system, if you use one. You must assign a dummy, or null, security profile to it.

Topology

This configuration relies on the deployment shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System" on page 57.

Configuration

IN THIS SECTION

- Configuring Logical System Security Profiles | 79

CLI Quick Configuration

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

```
set system security-profile master-profile policy maximum 65
set system security-profile master-profile policy reserved 60
set system security-profile master-profile zone maximum 22
set system security-profile master-profile zone reserved 17
set system security-profile master-profile flow-session maximum 3000
set system security-profile master-profile flow-session reserved 2100
set system security-profile master-profile icap-redirect-profile maximum 64
 set system security-profile master-profile icap-redirect-profile reserved 30
set system security-profile master-profile nat-nopat-address maximum 115
set system security-profile master-profile nat-nopat-address reserved 100
set system security-profile master-profile nat-static-rule maximum 125
set system security-profile master-profile nat-static-rule reserved 100
set system security-profile master-profile idp
set system security-profile master-profile root-logical-system
set system security-profile ls-accnt-mrkt-profile policy maximum 65
set system security-profile ls-accnt-mrkt-profile policy reserved 60
set system security-profile ls-accnt-mrkt-profile zone maximum 22
set system security-profile ls-accnt-mrkt-profile zone reserved 17
set system security-profile ls-accnt-mrkt-profile flow-session maximum 2500
set system security-profile ls-accnt-mrkt-profile flow-session reserved 2000
set system security-profile master-profile icap-redirect-profile maximum 64
 set system security-profile master-profile icap-redirect-profile reserved 30
set system security-profile ls-accnt-mrkt-profile nat-nopat-address maximum 125
set system security-profile ls-accnt-mrkt-profile nat-nopat-address reserved 100
set system security-profile ls-accnt-mrkt-profile nat-static-rule maximum 125
set system security-profile ls-accnt-mrkt-profile nat-static-rule reserved 100
set system security-profile ls-accnt-mrkt-profile logical-system ls-marketing-dept
set system security-profile ls-accnt-mrkt-profile logical-system ls-accounting-dept
set system security-profile ls-design-profile policy maximum 50
set system security-profile ls-design-profile policy reserved 40
set system security-profile ls-design-profile zone maximum 10
 set system security-profile ls-design-profile zone reserved 5
set system security-profile ls-design-profile flow-session maximum 2500
set system security-profile ls-design-profile flow-session reserved 2000
```

```
set system security-profile master-profile icap-redirect-profile maximum 64
 set system security-profile master-profile icap-redirect-profile reserved 30
 set system security-profile ls-design-profile nat-nopat-address maximum 120
set system security-profile ls-design-profile nat-nopat-address reserved 100
set system security-profile ls-design-profile logical-system ls-product-design
set system security-profile interconnect-profile logical-system interconnect-logical-system
```

Step-by-Step Procedure

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

Create three security profiles.

1. Create the first security profile.

   #### Step-by-Step Procedure

   a. Specify the number of maximum and reserved policies.

   ```
   [edit system security-profile]
   user@host# set master-profile policy maximum 65 reserved 60
   ```

   b. Specify the number of maximum and reserved zones.

   ```
   [edit system security-profile]
   user@host# set master-profile zone maximum 22 reserved 17
   ```

   c. Specify the number of maximum and reserved sessions.

   ```
   [edit system security-profile]
   user@host# set master-profile flow-session maximum 3000 reserved 2100
   ```

   d. Specify the number of maximum and reserved ICAP redirect profiles

   ```
   [edit system security-profile]
   user@host# set master-profile icap-redirect-profile maximum 64 reserved 30
   ```

e. Specify the number of maximum and reserved source NAT no-PAT addresses and static NAT rules.

```
[edit system security-profile]
user@host# set master-profile nat-nopat-address maximum 115 reserved 100
user@host# set master-profile nat-static-rule maximum 125 reserved 100
```

f. Enable intrusion detection and prevention (IDP). You can enable IDP only for the primary (root) logical system.

```
[edit system security-profile]
user@host# set idp
```

g. Bind the security profile to the logical system.

```
[edit system security-profile]
user@host# set master-profile root-logical-system
```

2. Create the second security profile.

Step-by-Step Procedure

a. Specify the number of maximum and reserved policies.

```
[edit system security-profile]
user@host# set ls-accnt-mrkt-profile policy maximum 65 reserved 60
```

b. Specify the number of maximum and reserved zones.

```
[edit system security-profile]
user@host# set ls-accnt-mrkt-profile zone maximum 22 reserved 17
```

c. Specify the number of maximum and reserved sessions.

```
[edit system security-profile]
user@host# set ls-accnt-mrkt-profile flow-session maximum 2500 reserved 2000
```

d. Specify the number of maximum and reserved ICAP redirect profiles

```
[edit system security-profile]
user@host# set ls-accnt-mrkt-profile icap-redirect-profile maximum 64 reserved 30
```

e. Specify the number of maximum and reserved source NAT no-PAT addresses.

```
[edit system security-profile]
user@host# set ls-accnt-mrkt-profile nat-nopat-address maximum 125 reserved 100
```

f. Specify the number of maximum and reserved static NAT rules.

```
[edit system security-profile]
user@host# set ls-accnt-mrkt-profile nat-static-rule maximum 125 reserved 100
```

g. Bind the security profile to two logical systems.

```
[edit system]
user@host# set security-profile ls-accnt-mrkt-profile logical-system ls-marketing-dept
user@host# set security-profile ls-accnt-mrkt-profile logical-system ls-accounting-dept
```

3. Create the third security profile.

Step-by-Step Procedure

a. Specify the number of maximum and reserved policies.

```
[edit system security-profile]
user@host# set ls-design-profile policy maximum 50 reserved 40
```

b. Specify the number of maximum and reserved zones.

```
[edit system security-profile]
user@host# set ls-design-profile zone maximum 10 reserved 5
```

c. Specify the number of maximum and reserved sessions.

```
[edit system security-profile]
user@host# set ls-design-profile flow-session maximum 2500 reserved 2000
```

d. Specify the number of maximum and reserved ICAP redirect profiles

```
[edit system security-profile]
user@host# setls-design-profile icap-redirect-profile maximum 64 reserved 30
```

e. Specify the number of maximum and reserved source NAT no-PAT addresses.

```
[edit system security-profile]
user@host# set ls-design-profile nat-nopat-address  maximum 120 reserved 100
```

4. Bind the security profile to a logical system.

```
user@host# set system security-profile ls-design-profile logical-system ls-product-design
```

5. Bind a null security profile to the interconnect logical system.

```
user@host# set system security-profile interconnect-profile logical-system interconnect-
logical-system
```

Results

From configuration mode, confirm your configuration by entering the `show system security-profile` command to see all security profiles configured.

To see individual security profiles, enter the `show system security-profile master-profile`, the `show system security-profile ls-accnt-mrkt-profile` and, the `show system security-profile ls-design-profile` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show system security-profile
    interconnect-profile {
        logical-system interconnect-logical-system;
```

```
    }
        ls-accnt-mrkt-profile {
            policy {
                maximum 65;
                reserved 60;
            }
            zone {
                maximum 22;
                reserved 17;
            }
            flow-session {
                maximum 2500;
                reserved 2000;
            }
            icap-redirect-profile {
                maximum 64;
                reserved 30;
            }
            nat-nopat-address {
                maximum 125;
                reserved 100;
            }
            nat-static-rule {
                maximum 125;
                reserved 100;
            }
            logical-system [ ls-marketing-dept ls-accounting-dept ];
        }
        ls-design-profile {
            policy {
                maximum 50;
                reserved 40;
            }
            zone {
                maximum 10;
                reserved 5;
            }
            flow-session {
                maximum 2500;
                reserved 2000;
            }
            icap-redirect-profile {
                maximum 64;
```

```
                reserved 30;
            }
            nat-nopat-address {
                maximum 120;
                reserved 100;
            }
            nat-static-rule {
                maximum 125;
                reserved 100;
            }
            logical-system ls-product-design;
        }
        master-profile {
            policy {
                maximum 65;
                reserved 60;
            }
            zone {
                maximum 22;
                reserved 17;
            }
            flow-session {
                maximum 3000;
                reserved 2100;
            }
            icap-redirect-profile {
                maximum 64;
                reserved 30;
            }
            nat-nopat-address {
                maximum 115;
                reserved 100;
            }
            nat-static-rule {
                maximum 125;
                reserved 100;
            }
```

```
            root-logical-system;
        }
```

```
user@host# show system security-profile master-profile
        policy {
            maximum 65;
            reserved 60;
            }
            zone {
                maximum 22;
                reserved 17;
            }
            flow-session {
                maximum 3000;
                reserved 2100;
            }
            icap-redirect-profile {
                maximum 64;
                reserved 30;
            }
            nat-nopat-address {
                maximum 115;
                reserved 100;
            }
            nat-static-rule {
                maximum 125;
                reserved 100;
            }
        root-logical-system;
```

```
user@host# show system security-profile ls-accnt-mrkt-profile
    policy {
        maximum 65;
        reserved 60;
        }
        zone {
            maximum 22;
            reserved 17;
        }
        flow-session {
```

```
            maximum 2500;
            reserved 2000;
        }
        icap-redirect-profile {
            maximum 64;
            reserved 30;
        }
        nat-nopat-address {
            maximum 125;
            reserved 100;
        }
        nat-static-rule {
            maximum 125;
            reserved 100;
        }
    logical-system [ ls-accounting-dept ls-marketing-dept ];
```

```
user@host# show system security-profile ls-design-profile
    policy {
        maximum 50;
        reserved 40;
        }
        zone {
            maximum 10;
            reserved 5;
        }
        flow-session {
            maximum 2500;
            reserved 2000;
        }
        icap-redirect-profile {
            maximum 64;
            reserved 30;
        }
        nat-nopat-address {
            maximum 120;
            reserved 100;
        }
        nat-static-rule {
            maximum 125;
            reserved 100;
```

```
    }
logical-system ls-product-design;
```

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

Verification

IN THIS SECTION

- Verifying That Security Profile Resources Are Effectively Allocated for Logical Systems | 88

To confirm that the security resources that you allocated for logical systems have been assigned to them, follow this procedure for each logical system and for all its resources.

Verifying That Security Profile Resources Are Effectively Allocated for Logical Systems

Purpose

Verify security resources for each logical system. Follow this process for all configured logical systems.

Action

1. Use SSH to log in to each user logical system as its user logical system administrator.

   Run SSH, specifying the IP address of your SRX Series device.

2. Enter the login ID and password for one of the user logical systems that you created.

```
login: lsmarketingadmin1
password: Talk2345
 lsmarketingadmin1@host:ls-marketing-dept>
```

3. Enter the following statement to identify the resources configured for the profile.

```
lsmarketingadmin1@host:ls-marketing-dept> show system security-profile ?
```

4. Enter the following command at the resulting prompt. Do this for each feature configured for the profile.

```
 lsmarketingadmin1@host:ls-marketing-dept> show system security-profile zone detail
logical system name : ls-marketing-dept
security profile name   : ls-accnt-mrkt-profile
 used amount             : 0
reserved amount         : 17
maximum quota           : 22
```

SEE ALSO

Understanding Logical Systems Security Profiles (Primary Administrators Only) | 71

Understanding the Primary Logical Systems and the Primary Administrator Role | 23

Understanding User Logical Systems and the User Logical System Administrator Role | 53

Example: Configuring User Logical Systems Security Profiles

IN THIS SECTION

- Requirements | 90
- Overview | 90
- Configuration | 92
- Verification | 95

In this example, you configure the user logical systems security profiles. It provides the information about a resource allocated to the logical system in a security profile.

NOTE:

- SRX4100 and SRX4200 devices support logical system in both transparent and route mode.

- SRX4600 device supports logical system in route mode only.
- Layer 2 cross logical system traffic is not supported.

Requirements

This example uses an SRX4100 and SRX4200 devices running Junos OS with logical systems.

Before you begin:

- Understand the logical system configuration process. See "User Logical Systems Configuration Overview" on page 51 to understand how this task fits into the overall configuration process.

Overview

Logical systems allow a primary administrator to partition an SRX Series device into discrete contexts called user logical systems. User logical systems are self-contained, private contexts, separate both from one another and from the primary logical system. A user logical system has its own security, networking, logical interfaces, routing configurations, and one or more user logical system administrators.

In this example, you configure security features for the user logical system described in Table 4 on page 91. This configuration used by the user logical system administrator to display resource information for a user logical system.

Table 4: Resource Information for a User Logical System

Field Name	Field Description
MAC flags	Status of MAC address learning properties for each interface: • **S**—Static MAC address is configured • **D**—Dynamic MAC address is configured • **L**—Locally learned MAC address is configured • **P**—Persistent static • **C**—Control MAC • **SE**—MAC accounting is enabled • **NM**—Non-configured MAC • **R**—Locally learned MAC address is configured • **O**—Open vSwitch Database (OVSDB) MAC
Ethernet switching table	For learned entries, the time at which the entry was added to the Ethernet switching table.
Logical system	Name of the logical system
Routing instance	Name of the routing instance
VLAN name	Name of the VLAN
MAC address	MAC address or addresses learned on a logical interface
Age	This field is not supported
Logical interface	Name of the logical interface

Table 4: Resource Information for a User Logical System *(Continued)*

Field Name	Field Description
RTR ID	ID of the routing device
NH Index	Software index of the next hop that is used to route the traffic for a given prefix.

Configuration

IN THIS SECTION

- CLI Quick Configuration | 92
- Procedure | 93
- Results | 95

CLI Quick Configuration

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

```
set system security-profile security-profile-name logical-system logical-system-name
set logical-systems logical-system-name interfaces xe-0/0/0 unit 0 family ethernet-switching
interface-mode access
set logical-systems logical-system-name interfaces xe-0/0/0 unit 0 family ethernet-switching
vlan members VLAN100
set logical-systems logical-system-name interfaces xe-0/0/1 unit 0 family ethernet-switching
interface-mode access
set logical-systems logical-system-name interfaces xe-0/0/1 unit 0 family ethernet-switching
vlan members VLAN100
set logical-systems logical-system-name interfaces xe-0/0/2 unit 0 family ethernet-switching
interface-mode trunk
set logical-systems logical-system-name interfaces xe-0/0/2 unit 0 family ethernet-switching
vlan members VLAN200
set logical-systems logical-system-name interfaces xe-0/0/1.0 unit 0 family ethernet-switching
```

```
interface-mode trunk
set logical-systems logical-system-name interfaces xe-0/0/2.0 unit 0 family ethernet-switching
vlan members vlan200
set logical-systems logical-system-name interfaces irb unit 22 family inet address
10.11.11.150/24
set logical-systems logical-system-name security policies default-policy permit-all
set logical-systems logical-system-name security zones security-zone trust host-inbound-traffic
system-services all
set logical-systems logical-system-name security zones security-zone trust host-inbound-traffic
protocols all
set logical-systems logical-system-name security zones security-zone trust interfaces xe-0/0/2.0
set logical-systems logical-system-name security zones security-zone untrust host-inbound-
traffic system-services all
set logical-systems logical-system-name security zones security-zone untrust host-inbound-
traffic protocols all
set logical-systems logical-system-name security zones security-zone untrust interfaces
xe-0/0/2.0
set logical-systems logical-system-name security zones security-zone untrust interfaces
xe-0/0/3.0
set logical-systems logical-system-name vlans VLAN100 vlan-id 100
set logical-systems logical-system-name vlans VLAN100 l3-interface irb.22
```

Procedure

Step-by-Step Procedure

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

To configure user logical systems security profiles:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

```
[edit]
admin@host> configure
admin@host#
```

2. Configure a security profile and assign it to a logical-system.

```
[edit system security-profile ]
admin@host# set system security-profile security-profile-name logical-system
```

3. Set the interfaces to the appropriate interface modes and specify that the logical interface that will receive the untagged data packets is a member of the native VLAN.

```
[edit logical-systems]
admin@host#set logical-systems logical-system-name interfaces xe-0/0/0 unit 0 family ethernet-switching interface-mode access
admin@host# set logical-systems logical-system-name interfaces xe-0/0/2 unit 0 family ethernet-switching vlan members VLAN100
admin@host#set logical-systems logical-system-name interfaces xe-0/0/1 unit 0 family ethernet-switching interface-mode access
admin@host# set logical-systems logical-system-name interfaces xe-0/0/3 unit 0 family ethernet-switching vlan members VLAN100
admin@host#set logical-systems logical-system-name interfaces xe-0/0/2 unit 0 family ethernet-switching interface-mode trunk
admin@host#set logical-systems logical-system-name interfaces xe-0/0/2 unit 0 family ethernet-switching vlan members VLAN100
admin@host#set logical-systems logical-system-name interfaces xe-0/0/1.0 unit 0 family ethernet-switching interface-mode trunk
admin@host#set logical-systems logical-system-name interfaces xe-0/0/2.0 unit 0 family ethernet-switching vlan members vlan200
```

4. Create the IRB interface and assign it an address in the subnet.

```
[edit interface]
admin@host# set interfaces irb unit 22 family inet address 10.11.11.150/24
```

5. Create the security policy to permit traffic from the trust zone to the untrust zone and assign interfaces to each zone.

```
[edit security policies]
admin@host# set security policies default-policy permit-all
admin@host# set security zones security-zone trust host-inbound-traffic system-services all
admin@host# set security zones security-zone trust host-inbound-traffic protocols all
admin@host# set security zones security-zone trust interfaces xe-0/0/2.0
```

```
admin@host# set security zones security-zone untrust host-inbound-traffic system-services all
admin@host# set security zones security-zone untrust host-inbound-traffic protocols all
admin@host# set security zones security-zone untrust interfaces xe-0/0/2.0
admin@host# set security zones security-zone untrust interfaces xe-0/0/3.0
```

6. Associate an IRB interface with the VLAN.

```
[edit logical-systems]
admin@host# set logical-systems logical-system-name vlans VLAN100 vlan-id 100
admin@host# set logical-systems logical-system-name vlans VLAN100 l3-interface irb.22
```

Results

From configuration mode, confirm your configuration by entering the `show ethernet-switching table` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
admin@host# show ethernet-switching table
ethernet-switching table {
    filter;
    inner-vlan;
    inter-switch-link;
    interface-mode;
    policer;
    recovery-timeout;
    storm-control;
    vlan;
    vlan-auto-sense;
    vlan-rewrite;
}
```

Verification

IN THIS SECTION

- Verifying User Logical Systems Security Profiles Configuration | 96

To confirm that the configuration is working properly, perform these tasks:

Verifying User Logical Systems Security Profiles Configuration

Purpose

Verify security policies information.

Action

From operational mode, enter the `show ethernet-switching table` command.

```
admin@host> show ethernet-switching table
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static, C -
Control MAC
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)


Ethernet switching table : 1 entries, 1 learned
Logical system   : LD2
Routing instance : default
    Vlan                MAC                 MAC         Age    Logical                NH
RTR
    name                address             flags              interface              Index
ID
    VLAN100             d4:04:ff:89:fd:30   D             -    xe-0/0/2.0             0
0
```

Example: Configuring Security log stream for Logical Systems

IN THIS SECTION

- Requirements | 97
- Overview | 97
- Configuration | 97
- Verification | 98

This example shows how to configure a security profiles for a logical system.

Requirements

This example uses the SRX Series devices running Junos OS with logical systems.

Before you begin:

- Read "SRX Series Logical Systems Primary Administrator Configuration Tasks Overview" on page 24 to understand how this task fits into the overall configuration process.
- See "Example: Configuring Logical Systems Security Profiles (Primary Administrators Only)" on page 77.

Overview

As primary administrator, you can configure a single security profile to assign resources to a specific logical system. Yo can use the same security profile for more than one logical system, or use a mix of both methods. The `set logical-system LSYS1 security log` command is introduced for logging support on SRX Series devices.

Configuration

IN THIS SECTION

- Configuring Logical System Security Profiles logical-system | 97
- Results | 98

Configuring Logical System Security Profiles logical-system

CLI Quick Configuration

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

```
set system security-profile p1 security-log-stream-number reserved 1
set system security-profile p1 security-log-stream-number maximum 2
set system security-profile p1 logical-system LSYS1
```

Step-by-Step Procedure

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

1. Configure a security profile and specify the number of maximum and reserved policies..

```
[edit system]
user@host# set security-profile p1 security-log-stream-number reserved 1
user@host# set security-profile p1 security-log-stream-number maximum 2
```

2. Assign the configured security profile to LSYS1.

```
user@host# set security-profile p1 logical-system LSYS1
```

Results

From configuration mode, confirm your configuration by entering the `show system security-profile` command to see all security profiles configured.

```
[edit]
user@host# show system security-profile
p1 {
    security-log-stream-number {
        maximum 2;
        reserved 1;
    }
    logical-system LSYS1;
}
```

Verification

IN THIS SECTION

- Verifying Security Profile Resources for Logical Systems | 99
- Verifying security-log-stream-number for logical-systems | 100

Verifying security-log-stream-number summary for logical-systems | 101
Verifying security-log-stream-number detail for logical-systems | 101

To confirm that the configuration is working properly, perform the below tasks:

Verifying Security Profile Resources for Logical Systems

Purpose

Verify the security resources for each logical system.

Action

From operational mode, enter the `show system security-profile all-resource`, `show system security-profile security-log-stream-number logical-system all`, `show system security-profile security-log-stream-number summary`, or `show system security-profile security-log-stream-number detail logical-system all` command to see the output:

show system security-profile all-resource

```
user@host> show system security-profile all-resource

resource                                usage    reserved     maximum
[logical system name:   root-logical-system]
[security profile name: Default-Profile]
address-book                                0           0         512
auth-entry                                  0           0  2147483647
cpu on CP                                0.00%       1.00%      80.00%
cpu on SPU                               0.00%       1.00%      80.00%
flow-gate                                   0           0      524288
flow-session                                2           0     6291456
nat-cone-binding                            0           0       65536
nat-destination-pool                        0           0        4096
nat-destination-rule                        0           0        8192
nat-nopat-address                           0           0     1048576
nat-pat-address                             0           0        2048
nat-port-ol-ipnumber                        0           0           4
nat-rule-referenced-prefix                  0           0     1048576
nat-source-pool                             0           0        2048
```

```
nat-source-rule                          0          0        8192
nat-static-rule                          0          0       20480
policy                                   0          0       40000
policy-with-count                        0          0        1024
scheduler                                0          0          64
zone                                     0          0         512
```

Meaning

The sample outputs displays information about the resources allocated to the logical system in a security profile. For each resource specified, the number used by the logical system and the configured maximum and reserved values are displayed.

Verifying security-log-stream-number for logical-systems

Purpose

Verify the security-log-stream-number for each logical system.

Action

From operational mode, enter the `show system security-profile security-log-stream-number logical-system all` command to see the output:

show system security-profile security-log-stream-number logical-system all

```
user@host> show system security-profile security-log-stream-number logical-system all
logical system name   security profile name       usage    reserved     maximum
root-logical-system   Default-Profile                1           0           3
LSYS1                 sp1                            0           1           3
LSYS2                 sp2                            1           0           3
```

Meaning

The sample output displays the information about a resource allocated to the logical system in a security profile with security profile name. For each resource specified, the number used by the logical system and the configured maximum and reserved values are displayed.

Verifying security-log-stream-number summary for logical-systems

Purpose

Verify the security-log-stream-number summary.

Action

From operational mode, enter the `show system security-profile security-log-stream-number summary` command to see the output:

show system security-profile security-log-stream-number summary

```
user@host> show system security-profile security-log-stream-number summary
global used amount      : 0
global maximum quota    : 32
global available amount : 32
total logical systems   : 1
total security profiles : 0
heaviest usage / user   : 0     / root-logical-system
lightest usage / user   : 0     / root-logical-system
```

Meaning

The sample output displays the summary information about the resource for all logical systems.

Verifying security-log-stream-number detail for logical-systems

Purpose

Verify the security-log-stream-number detail.

Action

From operational mode, enter the `show system security-profile security-log-stream-number detail logical-system all` command to see the output:

show system security-profile security-log-stream-number detail logical-system all

```
user@host> show system security-profile security-log-stream-number detail logical-system all
logical system name     : root-logical-system
security profile name   : Default-Profile
used amount             : 0
reserved amount         : 0
maximum quota           : 8
logical system name     : lsys0
security profile name   : lsys_profile
used amount             : 0
reserved amount         : 0
maximum quota           : 8
logical system name     : lsys1
security profile name   : lsys_profile
used amount             : 0
reserved amount         : 0
maximum quota           : 8
logical system name     : lsys2
security profile name   : lsys_profile
used amount             : 0
reserved amount         : 0
maximum quota           : 8
```

Meaning

The sample output displays the detailed level of output for all logical systems.

SEE ALSO

security-profile-resources | 876

CPU Allocation for Logical Systems

IN THIS SECTION

- Understanding CPU Allocation and Control | **103**
- Example: Configuring CPU Utilization (Primary Administrators Only) | **108**

The CPU allocation for logical systems assign the reserved CPU resources to a logical system used to calculate the amount of CPU usage based on the runtime utilization. For more information, see the following topics:

Understanding CPU Allocation and Control

IN THIS SECTION

- CPU Control | **104**
- Reserved CPU Utilization Quota for Logical Systems | **104**
- CPU Control Target | **105**
- Shared CPU Resources and CPU Quotas | **105**
- Monitoring CPU Utilization | **107**

When device CPU utilization is low, logical systems can acquire and use CPU resources above their allocated reserve quotas as long as the system-wide utilization remains within a stable range. CPU utilization on a device should never reach 100 percent because a device running at 100 percent CPU utilization might be slow to respond to management or system events or be unable to handle traffic bursts.

CPU resources are used on a first-come first-served basis. Without controls, logical systems can compete for CPU resources and drive CPU utilization up to 100 percent. You cannot rely on the configuration of static resources, such as security policies and zones, to directly control CPU usage because a logical system with small numbers of static resources allocated could still consume a large

amount of CPU. Instead, the primary administrator can enable CPU resource control and configure CPU utilization parameters for logical systems.

NOTE: Only the primary administrator can enable CPU control and configure CPU utilization parameters. User logical system administrators can use the `show system security-profile cpu` command to view CPU utilization for their logical systems.

This topic includes the following sections:

CPU Control

The primary administrator enables CPU control with the `cpu-control` *configuration statement* at the [`edit system security-profile resources`] hierarchy level.

NOTE: The `resources` security profile is a special security profile that contains global settings that apply to all logical systems in the device. Other security profiles configured by the primary administrator are bound to specific logical systems.

When CPU control is enabled, the primary administrator can then configure the following CPU utilization parameters:

- A reserved CPU quota is the percentage of CPU utilization that is guaranteed for a logical system.
- The CPU control target is the upper limit, in percent, for system-wide CPU utilization on the device under normal operating conditions.

Reserved CPU Utilization Quota for Logical Systems

A configured reserved CPU quota guarantees that a specified percentage of CPU is always available to a logical system. During runtime, CPU utilization by each logical system is measured every two seconds. The reserved CPU quota is used to calculate the amount of CPU each logical system can use based on the runtime utilization.

The primary administrator specifies the reserved CPU quota in a logical system security profile with the `cpu reserved` configuration statement at the [`edit system security-profile` *profile-name*] hierarchy level. The security profile is bound to one or more logical systems. Unlike other resources that are allocated to a logical system in a security profile, no maximum allowed quota can be configured for CPU utilization.

The Junos OS software checks to ensure that the sum of reserved CPU quotas for all logical systems on the device is less than 90 percent of the CPU control target value. If CPU control is enabled and reserved CPU quotas are not configured, the default reserved CPU quota for the primary logical system

is 1 percent and the default reserved CPU quota for user logical systems is 0 percent. The primary administrator can configure reserved CPU quotas even if CPU control is not enabled. The primary administrator can enable or disable CPU control without changing security profiles.

CAUTION: The primary 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.

CPU Control Target

CPU control target is the upper limit, in percent, for CPU utilization on the device under normal operating conditions. If CPU utilization on the device surpasses the configured target value, the Junos OS software initiates controls to bring CPU utilization between the target value and 90 percent of the target value. For example, if the CPU control target value is 80 and CPU utilization on the device surpasses 80 percent, then controls are initiated to bring CPU utilization within the range of 72 (90 percent of 80) and 80 percent.

During runtime, CPU utilization by each logical system is measured every two seconds. Dropping packets reduces the CPU usage for a logical system. If the CPU usage of a logical system exceeds its quota, CPU utilization control drops the packets received on that logical system. The packet drop rate is calculated every two seconds based on CPU utilization of all logical systems.

The primary administrator configures the CPU control target with the `cpu-control-target` configuration statement at the [`edit system security-profile resources`] hierarchy level. A stable level of CPU utilization should be relatively close to 100 percent but allow for bursts in CPU utilization. The primary administrator should configure the CPU control target level based on an understanding of the usage pattern of the logical system's deployment on the device.

CPU control must be enabled for the Junos OS software to control CPU usage. If the primary administrator enables CPU control without specifying a CPU control target value, the default CPU control target is 80 percent.

Shared CPU Resources and CPU Quotas

The sum of the reserved CPU quotas for all logical systems on the device must be less than 90 percent of the CPU control target; the difference is called the shared CPU resource. The shared CPU resource is dynamically allocated among the logical systems that need additional CPU. This means that a logical system can use more CPU than its reserved CPU quota.

The CPU quota for a logical system is the sum of its reserved CPU quota and its portion of the shared CPU resource. If multiple logical systems need more CPU resources, they split the shared CPU resource based on the relative weights of their reserved CPU quotas. Logical systems with larger reserved CPU quotas receive larger portions of the shared CPU resource. The goal for CPU control is to keep the

actual CPU utilization of a logical system at its CPU quota. If a logical system's CPU needs are greater than its CPU quota, packets are dropped for that logical system.

The following scenarios illustrate CPU control for logical systems. In each scenario, the CPU control target value is 80, which means that CPU controls will keep the maximum system-wide CPU utilization between 72 and 80 percent. The reserved CPU quotas for the logical systems are configured as follows: primary and lsys1 logical systems are 10 percent each and the lsys2 logical system is 5 percent.

CPU Utilization Scenario 1

In this scenario, each of the three logical systems needs 40 percent of CPU. Table 5 on page 106 shows the CPU quotas for each logical system. Because the CPU needed by each logical system is greater than its CPU quota, packets are dropped for each logical system.

Table 5: CPU Utilization Scenario 1

Logical System	Needed CPU	CPU Quotas	Packets Dropped?
primary	40%	28.8%	Yes
lsys1	40%	28.8%	Yes
lsys2	40%	14.4%	Yes

CPU Utilization Scenario 2

In this scenario, the primary logical system needs 25 percent of CPU while the two user logical systems need 40 percent. Table 6 on page 106 shows the CPU quota for the primary logical system is equal to the CPU it needs, so no packets are dropped for the primary logical system and CPU control monitors the CPU utilization of the primary logical system. Packets are dropped for lsys1 and lsys2.

Table 6: CPU Utilization Scenario 2

Logical System	Needed CPU	CPU Quotas	Packets Dropped?
primary	25%	25%	No
lsys1	40%	31.3%	Yes

Table 6: CPU Utilization Scenario 2 *(Continued)*

Logical System	Needed CPU	CPU Quotas	Packets Dropped?
lsys2	40%	15.6%	Yes

CPU Utilization Scenario 3

In this scenario, the primary and lsys2 logical systems need 5 percent and 3 percent of CPU, respectively, while lsys1 needs 40 percent. Table 7 on page 107 shows system-wide CPU utilization is 48 percent, which is less than 72 percent (90 percent of the CPU control target), so no packets are dropped and CPU control monitors all logical systems.

Table 7: CPU Utilization Scenario 3

Logical System	Needed CPU	CPU Quota	Packets Dropped?
primary	5%	5%	No
lsys1	40%	40%	No
lsys2	3%	3%	No

Monitoring CPU Utilization

CPU utilization can be monitored by either the primary administrator or the user logical system administrators. The primary administrator can monitor CPU utilization for the primary logical system, a specified user logical system, or all logical systems. User logical system administrators can only monitor CPU utilization for their logical system.

The `show system security-profile cpu` command shows the usage and drop rate in addition to the reserved CPU quota configured for the logical system. During runtime, CPU utilization by each logical system is measured every two seconds. The usage and drop rates displayed are the values at the interval prior to when the `show` command is run. If the `detail` option is not specified, the utilization of the central point (CP) and the average utilization of all services processing units (SPUs) is shown. The `detail` option displays the CPU utilization on each SPU.

The CPU utilization log file `lsys-cpu-utilization-log` contains utilization data for all logical systems on the device. Only the primary administrator can view the log file with the `show log lsys-cpu-utilization-log` command.

SEE ALSO

Example: Configuring CPU Utilization (Primary Administrators Only) | **108**

Understanding Logical Systems Security Profiles (Primary Administrators Only) | **71**

Example: Configuring CPU Utilization (Primary Administrators Only)

IN THIS SECTION

- Requirements | **108**
- Overview | **108**
- Configuration | **109**
- Verification | **111**

The primary administrator can enable CPU control and configure CPU utilization parameters. This example shows how to enable CPU utilization control and configure CPU utilization quotas and a control target.

Requirements

Before you begin:

- Log in to the primary logical system as the primary administrator. See "Understanding the Primary Logical Systems and the Primary Administrator Role" on page 23.
- Bind security profiles to the primary logical system and user logical systems configured on the device. See "Example: Configuring Logical Systems Security Profiles (Primary Administrators Only)" on page 77.

Overview

In this example, you enable CPU control and set the CPU control target to be 85 percent. You allocate reserved CPU quotas to the logical systems shown in "Example: Creating User Logical Systems, Their

Administrators, Their Users, and an Interconnect Logical System" on page 57. The logical systems are bound to the security profiles shown in Table 8 on page 109 and are assigned the reserved CPU quotas in the security profiles.

Table 8: Logical Systems, Security Profiles, and Reserved CPU Quotas

Logical System	Security Profile	Reserved CPU Quotas
root-logical-system (primary)	primary-profile	2 percent
ls-product-design	ls-design-profile	2 percent
ls-marketing-dept, ls-accounting-dept	ls-accnt-mrkt-profile	1 percent

Configuration

IN THIS SECTION

- Procedure | 109

Procedure

CLI Quick Configuration

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

```
set system security-profile resources cpu-control
set system security-profile resources cpu-control-target 85
set system security-profile master-profile cpu reserved 2
set system security-profile ls-design-profile cpu reserved 2
set system security-profile ls-accnt-mrkt-profile cpu reserved 1
```

Step-by-Step Procedure

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

To configure CPU utilization control parameters:

1. Log in to the primary logical system as the primary administrator and enter configuration mode.

```
[edit]
admin@host> configure
admin@host#
```

2. Enable CPU control.

```
[edit system security-profile resources]
admin@host# set cpu-control
```

3. Configure the CPU control target.

```
[edit system security-profile resources]
admin@host# set cpu-control-target 85
```

4. Configure the reserved CPU quotas in the security profiles.

```
[edit system]
admin@host# set security-profile security-profile master-profile cpu reserved 2
admin@host# set security-profile security-profile ls-design-profile cpu reserved 2
admin@host# set security-profile security-profile ls-accnt-mrkt-profile cpu reserved 1
```

Results

From configuration mode, confirm your configuration by entering the `show system security-profile` command. If the output does not display the intended configuration, repeat the \ instructions in this example to correct the configuration.

For brevity, this `show` command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
admin@host# show system security-profile
    resources {
        cpu-control;
        cpu-control-target 85;
    }
        ls-accnt-mrkt-profile {
            ...
            cpu {
                reserved 1;
            }
            logical-system [ ls-marketing-dept ls-accounting-dept ];
        }
        ls-design-profile {
            ...
            cpu {
                reserved 2;
            }
            logical-system ls-product-design;
        }
        master-profile {
            ...
            cpu {
                reserved 2;
            }
            logical-system root-logical-system;
        }
```

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

Verification

IN THIS SECTION

- Verifying CPU Utilization | 112

Confirm that the configuration is working properly.

Verifying CPU Utilization

Purpose

Display the configured reserved CPU quota, the actual CPU usage, and the drop rate.

Action

From operational mode, enter the `show system security-profile cpu logical-system all` command.

```
admin@host> show system security-profile cpu logical-system all
CPU control: TRUE
CPU control target: 85.00%
logical system name    profile name    CPU name    usage(%)    reserved(%)    drop rate(%)
root-logical-system    master-profile CP              0.10%          2.00%           0.00%
root-logical-system    master-Profile SPU             0.25%          2.00%           0.00%
ls-product-design      ls-design-profile CP           0.53%          2.00%           0.00%
ls-product-design      ls-design-profile SPU          0.26%          2.00%           0.00%
ls-marketing-dept      ls-acct-mrkt-profile CP        0.10%          1.00%           0.00%
ls-marketing-dept      ls-acct-mrkt-profile SPU       0.15%          1.00%           0.00%
ls-accounting-dept     ls-acct-mrkt-profile CP        0.23%          1.00%           0.00%
ls-accounting-dept     ls-acct-mrkt-profile SPU       0.34%          1.00%           0.00%
```

SEE ALSO

Understanding CPU Allocation and Control | 103

Understanding Logical Systems Security Profiles (Primary Administrators Only) | 71

Routing and Interfaces for Primary Logical Systems

IN THIS SECTION

- Understanding Logical Systems Interfaces and Routing Instances | 113
- Example: Configuring Interfaces, Routing Instances, and Static Routes for the Primary and Interconnect Logical Systems and Logical Tunnel Interfaces for the User Logical Systems (Primary Administrators Only) | 114
- Example: Configuring OSPF Routing Protocol for the Primary Logical Systems | 127

Logical systems enables you to configure the interfaces, routing instances and the routing protocol. The primary logical system administrator can display or clear the routing protocol parameters for all logical systems whereas the administrator for the user logical system can display or clear the protocol parameters for their own logical system. For more information, see the following topics:

Understanding Logical Systems Interfaces and Routing Instances

Logical interfaces on the device are allocated among the user logical systems by the primary administrator. The user logical system administrator configures the attributes of the interfaces, including IP addresses, and assigns them to routing instances and zones.

A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. There can be multiple routing tables for a single routing instance—for example, unicast IPv4, unicast IPv6, and multicast IPv4 routing tables can exist in a single routing instance. Routing protocol parameters and options control the information in the routing tables.

Interfaces and routing instances can be configured in the primary logical system and in user logical systems. Configuring an interface or routing instance in a logical system is the same as configuring an interface or routing instance on a device that is not configured for logical systems. Any routing instance created within a logical system is only applicable to that logical system.

The default routing instance, primary, refers to the main inet.0 routing table in the logical system. The primary routing instance is reserved and cannot be specified as a routing instance. Routes are installed in the primary routing instance by default, unless a routing instance is specified. Configure global routing options and protocols for the primary routing instance by including statements at the [**edit protocols**] and [**edit routing-options**] hierarchy levels in the logical system.

You can configure only virtual router routing instance type in a user logical system. Only one virtual private LAN service (VPLS) routing instance type can be configured on the device and it must be in the interconnect logical system.

The user logical system administrator can configure and view all attributes for an interface or routing instance in a user logical system. All attributes of an interface or routing instance in a user logical system are also visible to the primary administrator.

Multicast is a “one source, many destinations” method of traffic distribution, which means the destinations needing to receive the information from a particular source receive the traffic stream. The primary and user logical system administrators can configure a logical system to support multicast applications. The same multicast configurations to configure a device as a node in a multicast network can be used in a logical system.

SEE ALSO

- Example: Configuring Interfaces and Routing Instances for a User Logical Systems | **139**
- User Logical Systems Configuration Overview | **51**
- Understanding User Logical Systems and the User Logical System Administrator Role | **53**

Example: Configuring Interfaces, Routing Instances, and Static Routes for the Primary and Interconnect Logical Systems and Logical Tunnel Interfaces for the User Logical Systems (Primary Administrators Only)

IN THIS SECTION

- Requirements | **115**
- Overview | **115**
- Configuration | **118**
- Verification | **126**

This topic covers configuration of interfaces, static routes, and routing instances for the primary and interconnect logical systems. It also covers configuration of logical tunnel interfaces for user logical systems.

Requirements

The example uses an SRX5600 device running Junos operating system (Junos OS) with logical systems.

Before you begin:

- Read "SRX Series Logical Systems Primary Administrator Configuration Tasks Overview" on page 24 to understand how and where this procedure fits in the overall primary administrator configuration process.
- Read "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System" on page 57
- "Understanding the Interconnect Logical System and Logical Tunnel Interfaces" on page 11

Overview

IN THIS SECTION

Topology | 117

This scenario shows how to configure interfaces for the logical systems on the device, including an interconnect logical system.

- For the interconnect logical system, the example configures logical tunnel interfaces lt-0/0/0.0, lt-0/0/0.2, lt-0/0/0.4, and lt-0/0/0.6. The example configures a routing instance called vr-ic and assigns the interfaces to it.

  Because the interconnect logical system acts as a virtual switch, it is configured as a virtual private LAN service (VPLS) routing instance type. The interconnect logical system's lt-0/0/0 interfaces are configured with ethernet-vpls as the encapsulation type. The corresponding peer lt-0/0/0 interfaces in the primary and user logical systems are configured with Ethernet as the encapsulation type.

  - lt-0/0/0.0 connects to lt-0/0/0.1 on the root logical system.
  - lt-0/0/0.2 connects to lt-0/0/0.3 on the ls-product-design logical system.
  - lt-0/0/0.4 connects to lt-0/0/0.5 on the ls-marketing-dept logical system.
  - lt-0/0/0.6 connects to lt-0/0/0.7 on the ls-accounting-dept logical system.
- For the primary logical system, called root-logical-system, the example configures ge-0/0/4.0 and assigns it to the vr1-root routing instance. The example configures lt-0/0/0.1 to connect to

lt-0/0/0.0 on the interconnect logical system and assigns it to the vr1-root routing instance. The example configures static routes to allow for communication with other logical systems and assigns them to the vr1-root routing instance.

- For the ls-product-design logical system, the example configures lt-0/0/0.3 to connect to lt-0/0/0.2 on the interconnect logical system.
- For the ls-marketing-dept logical system, the example configures lt-0/0/0.5 to connect to lt-0/0/0.4 on the interconnect logical system.
- For the ls-accounting-dept logical system, the example configures lt-0/0/0.7 to connect to lt-0/0/0.6 on the interconnect logical system.

Figure 6 on page 117 shows the topology for this deployment including virtual routers and their interfaces for all logical systems.

Topology

Figure 6: Configuring Logical Tunnel Interfaces, Logical Interfaces, and Virtual Routers

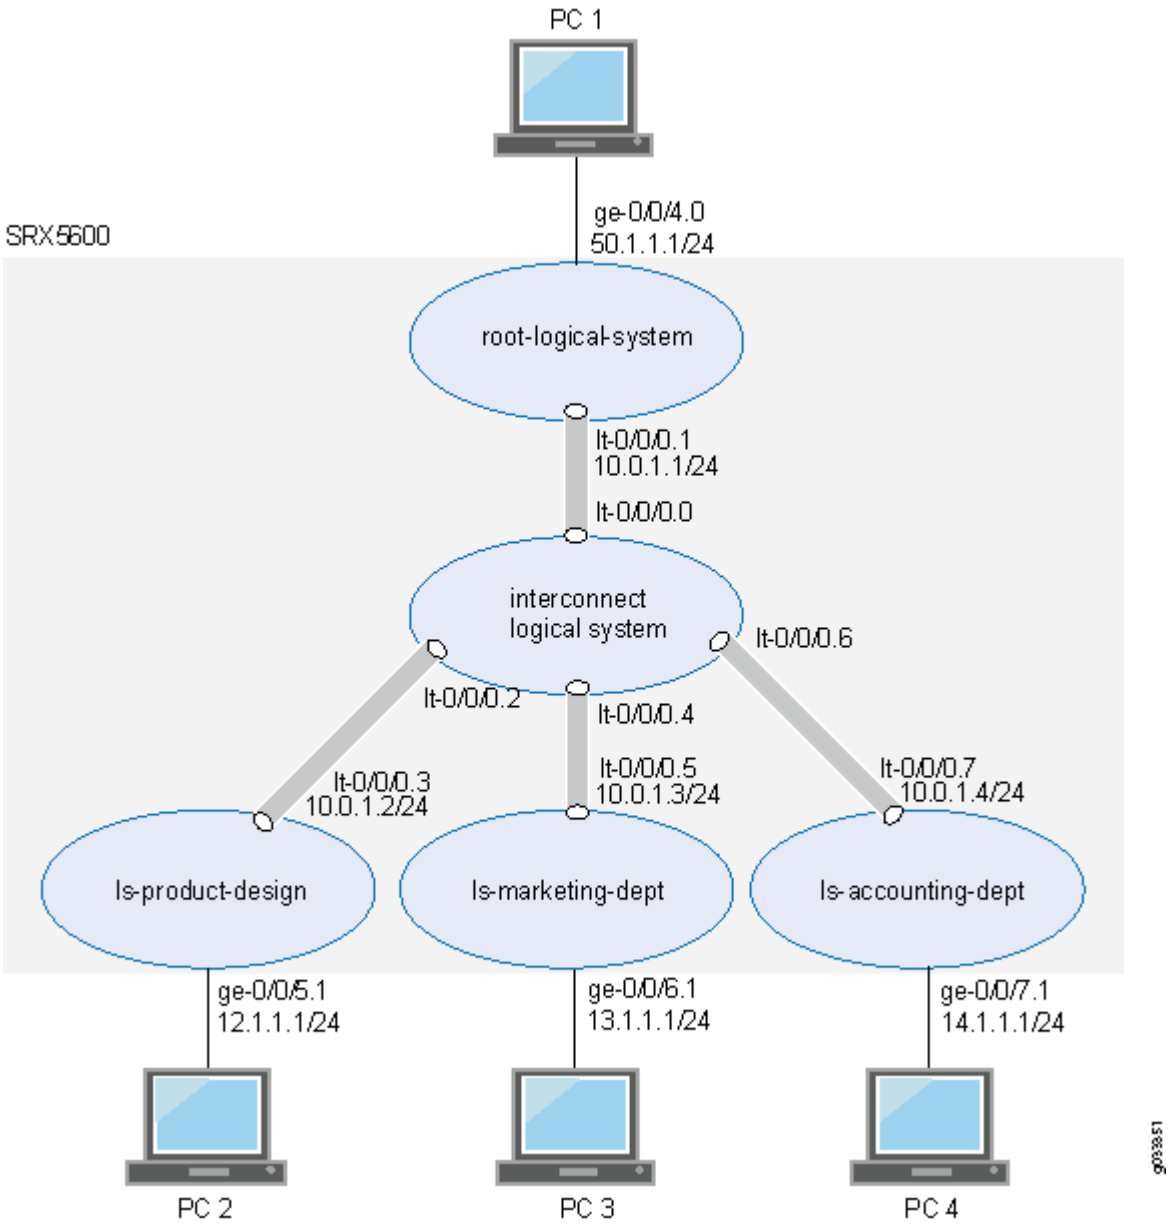

Configuration

IN THIS SECTION

- Configuring Logical Tunnel Interfaces and a Routing Instance for the Interconnect Logical System | 118
- Configuring Interfaces, a Routing Instance, and Static Routes for the Primary Logical System | 121
- Configuring Logical Tunnel Interfaces for the User Logical Systems | 124

This topic explains how to configure interfaces for logical systems.

Configuring Logical Tunnel Interfaces and a Routing Instance for the Interconnect Logical System

CLI Quick Configuration

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

```
set logical-systems interconnect-logical-system interfaces lt-0/0/0 unit 0 encapsulation
ethernet-vpls
set logical-systems interconnect-logical-system  interfaces lt-0/0/0 unit 0 peer-unit 1
set logical-systems interconnect-logical-system  interfaces lt-0/0/0 unit 2 encapsulation
ethernet-vpls
set logical-systems interconnect-logical-system  interfaces lt-0/0/0 unit 2 peer-unit 3
set logical-systems interconnect-logical-system  interfaces lt-0/0/0 unit 4 encapsulation
ethernet-vpls
set logical-systems interconnect-logical-system  interfaces lt-0/0/0 unit 4 peer-unit 5
set logical-systems interconnect-logical-system  interfaces lt-0/0/0 unit 6 encapsulation
ethernet-vpls
set logical-systems interconnect-logical-system  interfaces lt-0/0/0 unit 6 peer-unit 7
set logical-systems interconnect-logical-system routing-instances vr-ic instance-type vpls
set logical-systems interconnect-logical-system  routing-instances vr-ic interface lt-0/0/0.0
set logical-systems interconnect-logical-system  routing-instances vr-ic interface lt-0/0/0.2
set logical-systems interconnect-logical-system  routing-instances vr-ic interface lt-0/0/0.4
set logical-systems interconnect-logical-system  routing-instances vr-ic interface lt-0/0/0.6
```

Step-by-Step Procedure

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

To configure the interconnect system lt-0/0/0 interfaces and routing instances:

1. Configure the lt-0/0/0 interfaces.

```
[edit logical-systems]
user@host# set interconnect-logical-system interfaces lt-0/0/0 unit 0 encapsulation ethernet-vpls
user@host# set interconnect-logical-system  interfaces lt-0/0/0 unit 0 peer-unit 1
user@host# set interconnect-logical-system  interfaces lt-0/0/0 unit 2 encapsulation ethernet-vpls
user@host# set interconnect-logical-system  interfaces lt-0/0/0 unit 2 peer-unit 3
user@host# set interconnect-logical-system  interfaces lt-0/0/0 unit 4 encapsulation ethernet-vpls
user@host# set interconnect-logical-system  interfaces lt-0/0/0 unit 4 peer-unit 5
user@host# set interconnect-logical-system  interfaces lt-0/0/0 unit 6 encapsulation ethernet-vpls
user@host# set interconnect-logical-system  interfaces lt-0/0/0 unit 6 peer-unit 7
```

2. Configure the routing instance for the interconnect logical system and add its lt-0/0/0 interfaces to it.

```
[edit logical-systems]
user@host# set interconnect-logical-system routing-instances vr-ic  instance-type vpls
user@host# set interconnect-logical-system  routing-instances vr-ic  interface lt-0/0/0.0
user@host# set interconnect-logical-system  routing-instances vr-ic  interface lt-0/0/0.2
user@host# set interconnect-logical-system  routing-instances vr-ic  interface lt-0/0/0.4
user@host# set interconnect-logical-system  routing-instances vr-ic interface lt-0/0/0.6
```

Results

From configuration mode, confirm your configuration by entering the `show logical-systems interconnect-logical-system` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

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

```
user@host# show logical-systems interconnect-logical-system
    interfaces {
        lt-0/0/0 {
            unit 0 {
                encapsulation ethernet-vpls;
                peer-unit 1;
            }
            unit 2 {
                encapsulation ethernet-vpls;
                peer-unit 3;
            }
            unit 4 {
                encapsulation ethernet-vpls;
                peer-unit 5;
            }
            unit 6 {
                encapsulation ethernet-vpls;
                peer-unit 7;
            }
        }
    }
        routing-instances {
            vr-ic {
                instance-type vpls;
                interface lt-0/0/0.0;
                interface lt-0/0/0.2;
                interface lt-0/0/0.4;
                interface lt-0/0/0.6;
            }
        }
```

Configuring Interfaces, a Routing Instance, and Static Routes for the Primary Logical System

CLI Quick Configuration

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

```
set interfaces ge-0/0/4 vlan-tagging
set interfaces ge-0/0/4 unit 0 vlan-id 600
set interfaces ge-0/0/4 unit 0 family inet address 50.1.1.1/24
set interfaces ge-0/0/5 vlan-tagging
set interfaces ge-0/0/6 vlan-tagging
set interfaces ge-0/0/7 vlan-tagging
set interfaces lt-0/0/0 unit 1 encapsulation ethernet
set interfaces lt-0/0/0 unit 1 peer-unit 0
set interfaces lt-0/0/0 unit 1 family inet address 10.0.1.1/24
set routing-instances vr1-root instance-type virtual-router
set routing-instances vr1-root interface ge-0/0/4.0
set routing-instances vr1-root interface lt-0/0/0.1
set routing-instances vr1-root routing-options static route 12.1.1.0/24 next-hop 10.0.1.2
set routing-instances vr1-root routing-options static route 13.1.1.0/24 next-hop 10.0.1.3
set routing-instances vr1-root routing-options static route 14.1.1.0/24 next-hop 10.0.1.4
```

Step-by-Step Procedure

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

To configure the primary logical system interfaces:

1. Configure the primary (root) logical and lt-0/0/0.1 interfaces.

```
[edit interfaces]
user@host# set ge-0/0/4 vlan-tagging
user@host# set ge-0/0/4 unit 0 vlan-id 600
user@host# set ge-0/0/4 unit 0 family inet address 50.1.1.1/24
user@host# set lt-0/0/0 unit 1 encapsulation ethernet
user@host# set lt-0/0/0 unit 1 peer-unit 0
user@host# set lt-0/0/0 unit 1 family inet address 10.0.1.1/24
```

2. Configure the interfaces for other logical systems to support VLAN tagging.

```
[edit interfaces]
user@host# set ge-0/0/5 vlan-tagging
user@host# set ge-0/0/6 vlan-tagging
user@host# set ge-0/0/7 vlan-tagging
```

3. Configure a routing instance for the primary logical system, assign its interfaces to it, and configure static routes for it.

```
[edit routing-instances]
user@host# set vr1-root instance-type virtual-router
user@host# set vr1-root interface ge-0/0/4.0
user@host# set vr1-root interface lt-0/0/0.1
user@host# set vr1-root routing-options static route 12.1.1.0/24 next-hop 10.0.1.2
user@host# set vr1-root routing-options static route 13.1.1.0/24 next-hop 10.0.1.3
user@host# set vr1-root routing-options static route 14.1.1.0/24 next-hop 10.0.1.4
```

Results

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

```
[edit]
user@host# show interfaces
    ge-0/0/4 {
        vlan-tagging;
        unit 0 {
            vlan-id 600;
            family inet {
                address 50.1.1.1/24;
            }
        }
    }
        ge-0/0/5 {
            vlan-tagging;
        }
        ge-0/0/6 {
```

```
            vlan-tagging;
        }
        ge-0/0/7 {
            vlan-tagging;
        }
        lt-0/0/0 {
            unit 1 {
                encapsulation ethernet;
                peer-unit 0;
                family inet {
                    address 10.0.1.1/24;
                }
            }
        }
```

```
[edit]
user@host# show routing-instances
    vr1-root {
        instance-type virtual-router;
        interface ge-0/0/4.0;
        interface lt-0/0/0.1;
        routing-options {
            static {
                route 14.1.1.0/24 next-hop 10.0.1.4;
                route 12.1.1.0/24 next-hop 10.0.1.2;
                route 13.1.1.0/24 next-hop 10.0.1.3;
            }
        }
    }
```

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

Configuring Logical Tunnel Interfaces for the User Logical Systems

CLI Quick Configuration

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

```
set logical-systems ls-product-design interfaces lt-0/0/0 unit 3 encapsulation ethernet
set logical-systems ls-product-design interfaces lt-0/0/0 unit 3 peer-unit 2
set logical-systems ls-product-design interfaces lt-0/0/0 unit 3 family inet address 10.0.1.2/24
set logical-systems ls-marketing-dept interfaces lt-0/0/0 unit 5 encapsulation ethernet
set logical-systems ls-marketing-dept interfaces lt-0/0/0 unit 5 peer-unit 4
set logical-systems ls-marketing-dept interfaces lt-0/0/0 unit 5 family inet address 10.0.1.3/24
set logical-systems ls-accounting-dept interfaces lt-0/0/0 unit 7 encapsulation ethernet
set logical-systems ls-accounting-dept interfaces lt-0/0/0 unit 7 peer-unit 6
set logical-systems ls-accounting-dept interfaces lt-0/0/0 unit 7 family inet address
10.0.1.4/24
```

Step-by-Step Procedure

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

1. Configure the lt-0/0/0 interface for the first user logical system:

```
[edit logical-systems]
user@host# set ls-product-design interfaces lt-0/0/0 unit 3 encapsulation ethernet
user@host# set ls-product-design interfaces lt-0/0/0 unit 3 peer-unit 2
user@host# set ls-product-design interfaces lt-0/0/0 unit 3 family inet address 10.0.1.2/24
```

2. Configure the lt-0/0/0 interface for the second user logical system.

```
[edit logical-systems]
user@host# set ls-marketing-dept interfaces lt-0/0/0 unit 5 encapsulation ethernet
user@host# set ls-marketing-dept interfaces lt-0/0/0 unit 5 peer-unit 4
user@host# set ls-marketing-dept interfaces lt-0/0/0 unit 5 family inet address 10.0.1.3/24
face
```

3. Configure the lt-0/0/0 interface for the third user logical system.

```
[edit logical-systems]
user@host# set ls-accounting-dept interfaces lt-0/0/0 unit 7 encapsulation ethernet
user@host# set ls-accounting-dept interfaces lt-0/0/0 unit 7 peer-unit 6
user@host# set ls-accounting-dept interfaces lt-0/0/0 unit 7 family inet address 10.0.1.4/24
```

Results

From configuration mode, confirm your configuration by entering the `show logical-systems ls-product-design interfaces lt-0/0/0`, `show logical-systems ls-marketing-dept interfaces lt-0/0/0` , and `show logical-systems ls-accounting-dept interfaces lt-0/0/0` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host#  show logical-systems ls-product-design interfaces lt-0/0/0
```

```
lt-0/0/0 {
    unit 3 {
        encapsulation ethernet;
        peer-unit 2;
        family inet {
            address 10.0.1.2/24;
        }
    }
}
user@host# show logical-systems ls-marketing-dept interfaces lt-0/0/0
    lt-0/0/0 {
        unit 5 {
            encapsulation ethernet;
            peer-unit 4;
            family inet {
                address 10.0.1.3/24;
            }
        }
    }
}
user@host# show logical-systems ls-accounting-dept interfaces lt-0/0/0
lt-0/0/0 {
    unit 7 {
```

```
        encapsulation ethernet;
        peer-unit 6;
        family inet {
            address 10.0.1.4/24;
        }
    }
}
```

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

Verification

IN THIS SECTION

- Verifying That the Static Routes Configured for the Primary Administrator Are Correct | **126**

To confirm that the configuration is working properly, perform these tasks:

Verifying That the Static Routes Configured for the Primary Administrator Are Correct

Purpose

Verify if you can send data from the primary logical system to the other logical systems.

Action

From operational mode, use the `ping` command.

SEE ALSO

- Understanding the Primary Logical Systems and the Primary Administrator Role | **23**
- Understanding User Logical Systems and the User Logical System Administrator Role | **53**
- Understanding the Interconnect Logical System and Logical Tunnel Interfaces | **11**

Example: Configuring OSPF Routing Protocol for the Primary Logical Systems

IN THIS SECTION

- Requirements | 127
- Overview | 127
- Configuration | 128
- Verification | 130

This example shows how to configure OSPF for the primary logical system.

Requirements

Before you begin:

- Log in to the primary logical system as the primary administrator. See "Example: Configuring Root Password for Logical Systems" on page 55.
- Configure logical interfaces ge-0/0/4.0 and lt-0/0/0.1 for the primary logical system and assign them to the vr1-root routing instance. See "Example: Configuring Interfaces, Routing Instances, and Static Routes for the Primary and Interconnect Logical Systems and Logical Tunnel Interfaces for the User Logical Systems (Primary Administrators Only)" on page 114.

Overview

In this example, you configure OSPF for the primary logical system, called root-logical-system, shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System" on page 57.

This example enables OSPF routing on the ge-0/0/4.0 and lt-0/0/0.1 interfaces in the primary logical system. You configure the following routing policies to export routes from the Junos OS routing table into OSPF in the vr1-root routing instance:

- ospf-redist-direct—Routes learned from directly connected interfaces.
- ospf-redist-static—Static routes.
- ospf-to-ospf—Routes learned from OSPF.

Configuration

IN THIS SECTION

- Procedure | 128

Procedure

CLI Quick Configuration

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

```
set policy-options policy-statement ospf-redist-direct from protocol direct
set policy-options policy-statement ospf-redist-direct then accept
set policy-options policy-statement ospf-redist-static from protocol static
set policy-options policy-statement ospf-redist-static then accept
set policy-options policy-statement ospf-to-ospf from protocol ospf
set policy-options policy-statement ospf-to-ospf then accept
set routing-instances vr1-root protocols ospf export ospf-redist-direct
 set routing-instances vr1-root protocols ospf export ospf-redist-static
set routing-instances vr1-root protocols ospf export ospf-to-ospf
set routing-instances vr1-root protocols ospf area 0.0.0.1 interface ge-0/0/4.0
set routing-instances vr1-root protocols ospf area 0.0.0.1 interface lt-0/0/0.1
```

Step-by-Step Procedure

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

To configure OSPF for the primary logical system:

1. Log in to the primary logical system as the primary administrator and enter configuration mode.

```
admin@host> configure
admin@host#
```

2. Create routing policies that accept routes.

```
[edit policy-options]
admin@host# set policy-statement ospf-redist-direct from protocol direct
admin@host# set policy-statement ospf-redist-direct then accept
admin@host# set policy-statement ospf-redist-static from protocol static
admin@host# set policy-statement ospf-redist-static then accept
admin@host# set policy-statement ospf-to-ospf from protocol ospf
admin@host# set policy-statement ospf-to-ospf then accept
```

3. Apply the routing policies to routes exported from the Junos OS routing table into OSPF.

```
[edit routing-instances]
admin@host# set vr1-root protocols ospf export ospf-redist-direct
admin@host# set vr1-root protocols ospf export ospf-redist-static
admin@host# set vr1-root protocols ospf export ospf-to-ospf
```

4. Enable OSPF on the logical interfaces.

```
[edit routing-instances]
admin@host# set vr1-root protocols ospf area 0.0.0.1 interface ge-0/0/4.0
admin@host# set vr1-root protocols ospf area 0.0.0.1 interface lt-0/0/0.1
```

Results

From configuration mode, confirm your configuration by entering the `show policy-options` and `show routing-instances` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this `show` command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
    admin@host# show policy-options
    policy-statement ospf-redist-direct {
        from protocol direct;
        then accept;
    }
    policy-statement ospf-redist-static {
        from protocol static;
        then accept;
    }
    policy-statement ospf-to-ospf {
        from protocol ospf;
        then accept;
    }
    [edit]
admin@host# show routing-instances
    vr1-root {
        ...
        protocols {
            ospf {
                export [ ospf-redist-direct ospf-to-ospf ospf-redist-static ];
                area 0.0.0.1 {
                    interface lt-0/0/0.1;
                    interface ge-0/0/4.0;
                }
            }
        }
    }
```

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

Verification

IN THIS SECTION

- Verifying OSPF Interfaces | **131**
- Verifying OSPF Neighbors | **131**

- Verifying OSPF Routes | 132

Confirm that the configuration is working properly.

Verifying OSPF Interfaces

Purpose

Verify OSPF-enabled interfaces.

Action

From the CLI, enter the `show ospf interface instance vr1-root` command.

```
admin@host> show ospf interface instance vr1-root
Interface           State   Area            DR ID           BDR ID          Nbrs
lt-0/0/0.1          DR      0.0.0.0         10.0.1.1        0.0.0.0            0
ge-0/0/4.0          DR      0.0.0.1         10.0.1.1        0.0.0.0            0
```

Verifying OSPF Neighbors

Purpose

Verify OSPF neighbors.

Action

From the CLI, enter the `show ospf neighbor instance vr1-root` command.

```
admin@host> show ospf neighbor instance vr1-root
Address      Interface        State      ID           Pri    Dead
10.0.1.2     plt0.3         Full      0.0.0.0       128      39
```

Verifying OSPF Routes

Purpose

Verify OSPF routes.

Action

From the CLI, enter the `show ospf route instance vr1-root` command.

```
admin@host> show ospf route instance vr1-root
Topology default Route Table:

Prefix             Path  Route      NH       Metric NextHop       Nexthop
                   Type  Type       Type             Interface     Address/LSP
10.0.1.0/24        Intra Network    IP            1 lt-0/0/0.1
12.12.1.0/24       Intra Network    IP            1 ge-0/0/4.0
```

SEE ALSO

- Understanding Logical Systems Interfaces and Routing Instances | **113**
- Example: Configuring OSPF Routing Protocol for a User Logical Systems | **143**
- OSPF User Guide

Routing, Interfaces, and NAT for User Logical Systems

IN THIS SECTION

- Understanding Logical Systems Network Address Translation | **133**
- Example: Configuring Network Address Translation for a User Logical Systems | **134**
- Example: Configuring Interfaces and Routing Instances for a User Logical Systems | **139**

Example: Configuring OSPF Routing Protocol for a User Logical Systems | 143

The user logical system enables you to configure routing protocols, interfaces and NAT. Routing protocols handles all routing messages. NAT is a mechanism to translate the IP address of a computer or group of computers into a single public address when the packets are sent out to the internet. For more information, see the following topics:

Understanding Logical Systems Network Address Translation

Network Address Translation (NAT) is a method for modifying or translating network address information in packet headers. Either or both source and destination addresses in a packet may be translated. NAT can include the translation of port numbers as well as IP addresses.

Any combination of static, destination, or source NAT can be configured in the root or user logical systems. Configuring NAT in a logical system is the same as configuring NAT in a root system. The primary administrator can configure and monitor NAT in the primary logical system as well as any user logical system.

Starting in Junos OS Release 18.2R1, the NAT functionality is supported for logical systems on SRX4100, and SRX4200 devices in addition to existing support on SRX1500, SRX5400, SRX5600, and SRX5800 devices.

For each user logical system, the primary administrator can configure the maximum and reserved numbers for the following NAT resources:

- Source NAT pools and destination NAT pools
- IP addresses in source NAT pools with and without port address translation
- Rules for source, destination, and static NAT
- Persistent NAT bindings
- IP addresses that support port overloading

From a user logical system, the user logical system administrator can use the operational command `show system security-profile` with a NAT option to view the number of NAT resources allocated to the user logical system.

NOTE: The primary administrator can configure a security profile for the primary logical system that specifies the maximum and reserved numbers of NAT resources applied to the primary logical system. The number of resources configured in the primary logical system count toward the maximum number of NAT resources available on the device.

From a user logical system, the user logical system administrator can use the `show security nat` command to view the information about NAT for the user logical system. From the primary logical system, the primary administrator can use the same command to view information for the primary logical system, a specific user logical system, or all logical systems.

SEE ALSO

Example: Configuring Network Address Translation for a User Logical Systems | **134**

User Logical Systems Configuration Overview | **51**

Understanding Logical Systems Security Profiles (Primary Administrators Only) | **71**

Introduction to NAT

Example: Configuring Network Address Translation for a User Logical Systems

IN THIS SECTION

- Requirements | **134**
- Overview | **135**
- Configuration | **135**
- Verification | **138**

This example shows how to configure static NAT for a user logical system.

Requirements

Before you begin:

- Log in to the user logical system as the logical system administrator. See "User Logical Systems Configuration Overview" on page 51.
- Use the `show system security-profile nat-static-rule` command to see the static NAT resources allocated to the logical system.
- Configure security policies. See "Example: Configuring Security Policies in a User Logical Systems" on page 216.

Overview

This example configures the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System" on page 57.

Devices in the ls-product-design-untrust zone access a specific host in the ls-product-design-trust zone by way of the address 12.1.1.200/32. For packets that enter the ls-product-design logical system from the ls-product-design-untrust zone with the destination IP address 12.1.1.200/32, the destination IP address is translated to the 12.1.1.100/32. This example configures the static NAT described in Table 9 on page 135.

Table 9: User Logical System Static NAT Configuration

Feature	Name	Configuration Parameters
Static NAT rule set	rs1	• Rule r1 to match packets from the ls-product-design-untrust zone with destination address 12.1.1.200/32. • Destination IP address in matching packets is translated to 12.1.1.100/32.
Proxy ARP		Address 12.1.1.200 on interface lt-0/0/0.3.

Configuration

IN THIS SECTION

- Procedure | 136

Procedure

CLI Quick Configuration

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

```
set security nat static rule-set rs1 from zone ls-product-design-untrust
set security nat static rule-set rs1 rule r1 match destination-address 12.1.1.200/32
set security nat static rule-set rs1 rule r1 then static-nat prefix 12.1.1.100/32
set security nat proxy-arp interface lt-0/0/0.3 address 12.1.1.200/32
```

Step-by-Step Procedure

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

To configure NAT in a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

```
lsdesignadmin1@host:ls-product-design> configure
lsdesignadmin1@host:ls-product-design#
```

2. Configure a static NAT rule set.

```
[edit security nat static]
lsdesignadmin1@host:ls-product-design# set rule-set rs1 from zone ls-product-design-untrust
```

3. Configure a rule that matches packets and translates the destination address in the packets.

```
[edit security nat static]
lsdesignadmin1@host:ls-product-design# set rule-set rs1 rule r1 match destination-address
12.1.1.200/32
lsdesignadmin1@host:ls-product-design# set rule-set rs1 rule r1 then static-nat prefix
12.1.1.100/32
```

4. Configure proxy ARP.

```
[edit security nat]
lsdesignadmin1@host:ls-product-design# set proxy-arp interface lt-0/0/0.3 address
12.1.1.200/32
```

Results

From configuration mode, confirm your configuration by entering the `show security nat` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
lsdesignadmin1@host:ls-product-design# show security nat
static {
    rule-set rs1 {
        from zone ls-product-design-untrust;
        rule r1 {
            match {
                destination-address 12.1.1.200/32;
            }
            then {
                static-nat prefix 12.1.1.100/32;
            }
        }
    }
}
    proxy-arp {
        interface lt-0/0/0.3 {
            address {
                12.1.1.200/32;
            }
        }
    }
```

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

Verification

IN THIS SECTION

- Verifying Static NAT Configuration | 138
- Verifying NAT Application to Traffic | 138

To confirm that the configuration is working properly, perform these tasks:

Verifying Static NAT Configuration

Purpose

Verify that there is traffic matching the static NAT rule set.

Action

From operational mode, enter the `show security nat static rule` command. View the Translation hits field to check for traffic that matches the rule.

Verifying NAT Application to Traffic

Purpose

Verify that NAT is being applied to the specified traffic.

Action

From operational mode, enter the `show security flow session` command.

SEE ALSO

User Logical Systems Configuration Overview | 51

Understanding Logical Systems Network Address Translation | 133

Static NAT Configuration Overview

Example: Configuring Interfaces and Routing Instances for a User Logical Systems

IN THIS SECTION

- Requirements | 139
- Overview | 139
- Configuration | 140

This example shows how to configure interfaces and routing instances for a tenant system.

Requirements

Before you begin:

- Log in to the user logical system as the user logical system administrator. See "User Logical Systems Configuration Overview" on page 51.
- Determine which logical interfaces and, optionally, which logical tunnel interfaces are allocated to your user logical system by the primary administrator. The primary administrator configures the logical tunnel interfaces. See "Understanding the Primary Logical Systems and the Primary Administrator Role" on page 23.

Overview

This example configures the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System" on page 57.

This example configures the interfaces and routing instances described in Table 10 on page 139.

Table 10: User Logical System Interface and Routing Instance Configuration

Feature	Name	Configuration Parameters
Interface	ge-0/0/5.1	• IP address 12.1.1.1/24 • VLAN ID 700

Table 10: User Logical System Interface and Routing Instance Configuration *(Continued)*

Feature	Name	Configuration Parameters

Routing instance	pd-vr1	• Instance type: virtual router
		• Includes interfaces ge-0/0/5.1 and lt-0/0/0.3
		• Static routes:
		  • 13.1.1.0/24 next-hop 10.0.1.3
		  • 14.1.1.0/24 next-hop 10.0.1.4
		  • 12.12.1.0/24 next-hop 10.0.1.1

Configuration

IN THIS SECTION

- Procedure | 140

Procedure

CLI Quick Configuration

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

```
set interfaces ge-0/0/5 unit 1 family inet address 12.1.1.1/24
set interfaces ge-0/0/5 unit 1 vlan-id 700
set routing-instances pd-vr1 instance-type virtual-router
set routing-instances pd-vr1 interface ge-0/0/5.1
set routing-instances pd-vr1 interface lt-0/0/0.3
set routing-instances pd-vr1 routing-options static route 13.1.1.0/24 next-hop 10.0.1.3
set routing-instances pd-vr1 routing-options static route 14.1.1.0/24 next-hop 10.0.1.4
set routing-instances pd-vr1 routing-options static route 12.12.1.0/24 next-hop 10.0.1.1
```

Step-by-Step Procedure

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

To configure an interface and a routing instance in a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

```
lsdesignadmin1@host:ls-product-design> configure
lsdesignadmin1@host:ls-product-design#
```

2. Configure the logical interface for a user logical system.

```
[edit interfaces]
lsdesignadmin1@host:ls-product-design# set ge-0/0/5 unit 1 family inet address 12.1.1.1/24
lsdesignadmin1@host:ls-product-design# set ge-0/0/5 unit 1 vlan-id 700
```

3. Configure the routing instance and assign interfaces.

```
[edit routing-instances]
lsdesignadmin1@host:ls-product-design#  set pd-vr1 instance-type virtual-router
lsdesignadmin1@host:ls-product-design#  set pd-vr1 interface ge-0/0/5.1
lsdesignadmin1@host:ls-product-design#  set pd-vr1 interface lt-0/0/0.3
```

4. Configure static routes.

```
[edit routing-instances]
lsdesignadmin1@host:ls-product-design# set pd-vr1 routing-options static route 13.1.1.0/24
next-hop 10.0.1.3
lsdesignadmin1@host:ls-product-design# set pd-vr1 routing-options static route 14.1.1.0/24
next-hop 10.0.1.4
lsdesignadmin1@host:ls-product-design# set pd-vr1 routing-options static route 12.12.1.0/24
next-hop 10.0.1.1
```

Results

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

NOTE: The primary administrator configures the lt-0/0/0.3 interface. Thus, the lt-0/0/0.3 configuration appears in the `show interfaces` output even though you did not configure this item.

```
lsdesignadmin1@host:ls-product-design# show interfaces
ge-0/0/5 {
    unit 1 {
        vlan-id 700;
        family inet {
            address 12.1.1.1/24;
        }
    }
}
    lt-0/0/0 {
        unit 3 {
            encapsulation ethernet;
            peer-unit 2;
            family inet {
                address 10.0.1.2/24;
            }
        }
    }
lsdesignadmin1@host:ls-product-design# show routing-instances
pd-vr1 {
    instance-type virtual-router;
    interface ge-0/0/5.1;
    interface lt-0/0/0.3;
    routing-options {
        static {
            route 13.1.1.0/24 next-hop 10.0.1.3;
            route 14.1.1.0/24 next-hop 10.0.1.4;
            route 12.12.1.0/24 next-hop 10.0.1.1;
        }
    }
}
```

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

SEE ALSO

User Logical Systems Configuration Overview | 51

Understanding Logical Systems Interfaces and Routing Instances | 113

Example: Configuring OSPF Routing Protocol for a User Logical Systems

IN THIS SECTION

- Requirements | 143
- Overview | 143
- Configuration | 144
- Verification | 147

This example shows how to configure OSPF for a user logical system.

Requirements

Before you begin:

- Log in to the user logical system as the user logical system administrator. See "User Logical Systems Configuration Overview" on page 51.
- Configure logical interface ge-0/0/5.1. Assign ge-0/0/5.1 and lt-0/0/0.3 to the pd-vr1 routing instance. See "Example: Configuring Interfaces and Routing Instances for a User Logical Systems" on page 139.

Overview

In this example, you configure OSPF for the ls-product-design user logical system, shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System" on page 57.

This example enables OSPF routing on the ge-0/0/5.1 and lt-0/0/0.3 interfaces in the ls-product-design user logical system. You configure the following routing policies to export routes from the Junos OS routing table into OSPF in the pd-vr1 routing instance:

- ospf-redist-direct—Routes learned from directly connected interfaces.
- ospf-redist-static—Static routes.
- ospf-to-ospf—Routes learned from OSPF.

Configuration

IN THIS SECTION

- Procedure | 144

Procedure

CLI Quick Configuration

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

```
set policy-options policy-statement ospf-redist-direct from protocol direct
set policy-options policy-statement ospf-redist-direct then accept
set policy-options policy-statement ospf-redist-static from protocol static
set policy-options policy-statement ospf-redist-static then accept
set policy-options policy-statement ospf-to-ospf from protocol ospf
set policy-options policy-statement ospf-to-ospf then accept
set routing-instances pd-vr1 protocols ospf export ospf-redist-direct
 set routing-instances pd-vr1 protocols ospf export ospf-redist-static
set routing-instances pd-vr1 protocols ospf export ospf-to-ospf
set routing-instances pd-vr1 protocols ospf area 0.0.0.1 interface ge-0/0/5.1
set routing-instances pd-vr1 protocols ospf area 0.0.0.1 interface lt-0/0/0.3
```

Step-by-Step Procedure

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

To configure OSPF for the user logical system:

1. Log in to the user logical system as the user logical system administrator and enter configuration mode.

```
lsdesignadmin1@host:ls-product-design> configure
lsdesignadmin1@host:ls-product-design#
```

2. Create routing policies that accept routes.

```
[edit policy-options]
lsdesignadmin1@host:ls-product-design# set policy-statement ospf-redist-direct from protocol
direct
lsdesignadmin1@host:ls-product-design# set policy-statement ospf-redist-direct then accept
lsdesignadmin1@host:ls-product-design# set policy-statement ospf-redist-static from protocol
static
lsdesignadmin1@host:ls-product-design# set policy-statement ospf-redist-static then accept
lsdesignadmin1@host:ls-product-design# set policy-statement ospf-to-ospf from protocol ospf
lsdesignadmin1@host:ls-product-design# set policy-statement ospf-to-ospf then accept
```

3. Apply the routing policies to routes exported from the Junos OS routing table into OSPF.

```
[edit routing-instances]
lsdesignadmin1@host:ls-product-design# set pd-vr1 protocols ospf export ospf-redist-direct
lsdesignadmin1@host:ls-product-design# set pd-vr1 protocols ospf export ospf-redist-static
lsdesignadmin1@host:ls-product-design# set pd-vr1 protocols ospf export ospf-to-ospf
```

4. Enable OSPF on the logical interfaces.

```
[edit routing-instances]
lsdesignadmin1@host:ls-product-design# set pd-vr1 protocols ospf area 0.0.0.1 interface
ge-0/0/5.1
lsdesignadmin1@host:ls-product-design# set pd-vr1 protocols ospf area 0.0.0.1 interface
lt-0/0/0.3
```

Results

From configuration mode, confirm your configuration by entering the `show policy-options` and `show routing-instances` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this `show` command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
lsdesignadmin1@host:ls-product-design# show policy-options
    policy-statement ospf-redist-direct {
        from protocol direct;
        then accept;
    }
    policy-statement ospf-redist-static {
        from protocol static;
        then accept;
    }
    policy-statement ospf-to-ospf {
        from protocol ospf;
        then accept;
    }
[edit]
lsdesignadmin1@host:ls-product-design# show routing-instances
    pd-vr1 {
        ...
        protocols {
            ospf {
                export [ ospf-redist-direct ospf-to-ospf ospf-redist-static ];
                area 0.0.0.1 {
                    interface lt-0/0/0.3;
                    interface ge-0/0/5.1;
                }
            }
        }
    }
```

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

Verification

IN THIS SECTION

- Verifying OSPF Interfaces | 147
- Verifying OSPF Neighbors | 147
- Verifying OSPF Routes | 148

Confirm that the configuration is working properly.

Verifying OSPF Interfaces

Purpose

Verify OSPF-enabled interfaces.

Action

From the CLI, enter the `show ospf interface instance pd-vr1` command.

```
lsdesignadmin1@host:ls-product-design> show ospf interface instance pd-vr1
Interface           State   Area            DR ID           BDR ID          Nbrs
lt-0/0/0.3          DR      0.0.0.0         10.0.1.2        0.0.0.0            0
ge-0/0/5.1          DR      0.0.0.1         10.0.1.2        0.0.0.0            0
```

Verifying OSPF Neighbors

Purpose

Verify OSPF neighbors.

Action

From the CLI, enter the `show ospf neighbor instance pd-vr1` command.

```
lsdesignadmin1@host:ls-product-design> show ospf neighbor instance pd-vr1
Address      Interface        State      ID           Pri    Dead
10.0.1.1     plt0.1         Full      0.0.0.0       128      39
```

Verifying OSPF Routes

Purpose

Verify OSPF routes.

Action

From the CLI, enter the `show ospf route instance pd-vr1` command.

```
lsdesignadmin1@host:ls-product-design> show ospf route instance pd-vr1
Topology default Route Table:

Prefix             Path  Route      NH       Metric NextHop
Nexthop                             Type  Type       Type            Interface     Address/LSP
10.0.1.0/24        Intra Network    IP            1 lt-0/0/0.3
12.12.1.0/24       Intra Network    IP            1 ge-0/0/5.1
```

SEE ALSO

- Understanding Logical Systems Interfaces and Routing Instances | 113
- Example: Configuring OSPF Routing Protocol for the Primary Logical Systems | 127
- OSPF User Guide

RELATED DOCUMENTATION

- User Logical Systems Overview | 51

Security Zones in Logical Systems

IN THIS SECTION

- Understanding Logical Systems Zones | 149
- Example: Configuring User Logical Systems | 150
- Example: Configuring Security Zones for a User Logical Systems | 167

Security zones are the building blocks for policies. Security zones are logical entities to which one or more interfaces are bound and provides a means of distinguishing groups of hosts (user logical systems and other hosts, such as servers), resources from one another in order to apply different security measures. For more information, see the following topics:

Understanding Logical Systems Zones

Security zones are logical entities to which one or more interfaces are bound. Security zones can be configured on the primary logical system by the primary administrator or on user logical systems by the user logical system administrator. On a logical system, the administrator can configure multiple security zones, dividing the network into network segments to which various security options can be applied.

The primary administrator configures the maximum and reserved numbers of security zones for each user logical system. The user logical system administrator can then create security zones in the user logical system and assign interfaces to each security zone. From a user logical system, the user logical system administrator can use the `show system security-profile zones` command to view the number of security zones allocated to the user logical system and the `show interfaces` command to view the interfaces allocated to the user logical system.

NOTE: The primary administrator can configure a security profile for the primary logical system that specifies the maximum and reserved numbers of security zones applied to the primary logical system. The number of zones configured in the primary logical system count toward the maximum number of zones available on the device.

The primary and user administrator can configure the following properties of a security zone in a logical system:

- Interfaces that are part of a security zone.
- Screen options—For every security zone, you can enable a set of predefined screen options that detect and block various kinds of traffic that the device determines as potentially harmful.
- TCP-Reset—When this feature is enabled, the system sends a TCP segment with the RESET flag set when traffic arrives that does not match an existing session and does not have the synchronize flag set.
- Host inbound traffic—This feature specifies the kinds of traffic that can reach the device from systems that are directly connected to its interfaces. You can configure these parameters at the zone level, in which case they affect all interfaces of the zone, or at the interface level. (Interface configuration overrides that of the zone.)

There are no preconfigured security zones in the primary logical system or user logical system.

The management functional zone (MGT) can only be configured for the primary logical system. There is only one management interface per device and that interface is allocated to the primary logical system.

The `all` interface can only be assigned to a zone in the primary logical system by the primary administrator.

The user logical system administrator can configure and view all attributes for a security zone in a user logical system. All attributes of a security zone in a user logical system are also visible to the primary administrator.

SEE ALSO

Example: Configuring Security Zones for a User Logical Systems | **167**

User Logical Systems Configuration Overview | **51**

Understanding Logical Systems Security Profiles (Primary Administrators Only) | **71**

Understanding Logical Systems Interfaces and Routing Instances | **113**

Security Zones Overview

Example: Configuring User Logical Systems

IN THIS SECTION

Requirements | **151**

Overview | 151
Configuration | 155
Verification | 167

This example shows the configuration of interfaces, routing instances, zones, and security policies for user logical systems.

Requirements

Before you begin:

- Log in to the user logical system as the logical system administrator. See "User Logical Systems Configuration Overview" on page 51.
- Be sure you know which logical interfaces and optionally, which logical tunnel interface (and its IP address) are allocated to your user logical system by the primary administrator. See "Understanding the Primary Logical Systems and the Primary Administrator Role" on page 23.

Overview

This example configures the ls-marketing-dept and ls-accounting-dept user logical systems shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System" on page 57.

This example configures the parameters described in Table 11 on page 151 and Table 12 on page 153.

Table 11: ls-marketing-dept Logical System Configuration

Feature	Name	Configuration Parameters
Interface	ge-0/0/6.1	• IP address 13.1.1.1/24 • VLAN ID 800

Table 11: ls-marketing-dept Logical System Configuration *(Continued)*

Feature	Name	Configuration Parameters

Routing instance	mk-vr1	• Instance type: virtual router
		• Includes interfaces ge-0/0/6.1 and lt-0/0/0.5
		• Static routes:
		  • 12.1.1.0/24 next-hop 10.0.1.2
		  • 14.1.1.0/24 next-hop 10.0.1.4
		  • 12.12.1.0/24 next-hop 10.0.1.1

Zones	ls-marketing-trust	Bind to interface ge-0/0/6.1.
	ls-marketing-untrust	Bind to interface lt-0/0/0.5
Address books	marketing-internal	• Address marketers: 13.1.1.0/24 • Attach to zone ls-marketing-trust
	marketing-external	• Address design: 12.1.1.0/24 • Address accounting: 14.1.1.0/24 • Address others: 12.12.1.0/24 • Address set otherlsys: design, accounting • Attach to zone ls-marketing-untrust

Table 11: ls-marketing-dept Logical System Configuration *(Continued)*

Feature	Name	Configuration Parameters
Policies	permit-all-to-otherlsys	Permit the following traffic: • From zone: ls-marketing-trust • To zone: ls-marketing-untrust • Source address: marketers • Destination address: otherlsys • Application: any
	permit-all-from-otherlsys	Permit the following traffic: • From zone: ls-marketing-untrust • To zone: ls-marketing-trust • Source address: otherlsys • Destination address: marketers • Application: any

Table 12: ls-accounting-dept Logical System Configuration

Feature	Name	Configuration Parameters
Interface	ge-0/0/7.1	• IP address 14.1.1.1/24 • VLAN ID 900

Table 12: ls-accounting-dept Logical System Configuration *(Continued)*

Feature	Name	Configuration Parameters

Routing instance	acct-vr1	• Instance type: virtual router
		• Includes interfaces ge-0/0/7.1 and lt-0/0/0.7
		• Static routes:
		  • 12.1.1.0/24 next-hop 10.0.1.2
		  • 13.1.1.0/24 next-hop 10.0.1.3
		  • 12.12.1.0/24 next-hop 10.0.1.1

Zones	ls-accounting-trust	Bind to interface ge-0/0/7.1.
	ls-accounting-untrust	Bind to interface lt-0/0/0.7
Address books	accounting-internal	• Address accounting: 14.1.1.0/24 • Attach to zone ls-accounting-trust
	accounting-external	• Address design: 12.1.1.0/24 • Address marketing: 13.1.1.0/24 • Address others: 12.12.1.0/24 • Address set otherlsys: design, marketing • Attach to zone ls-accounting-untrust

Table 12: ls-accounting-dept Logical System Configuration *(Continued)*

Feature	Name	Configuration Parameters
Policies	permit-all-to-otherlsys	Permit the following traffic: • From zone: ls-accounting-trust • To zone: ls-accounting-untrust • Source address: accounting • Destination address: otherlsys • Application: any
	permit-all-from-otherlsys	Permit the following traffic: • From zone: ls-accounting-untrust • To zone: ls-accounting-trust • Source address: otherlsys • Destination address: accounting • Application: any

Configuration

IN THIS SECTION

- Configuring the ls-marketing-dept User Logical System | **156**
- Configuring the ls-accounting-dept User Logical System | **161**

Configuring the ls-marketing-dept User Logical System

CLI Quick Configuration

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

```
set interfaces ge-0/0/6 unit 1 family inet address 13.1.1.1/24
set interfaces ge-0/0/6 unit 1 vlan-id 800
set routing-instances mk-vr1 instance-type virtual-router
set routing-instances mk-vr1 interface ge-0/0/6.1
set routing-instances mk-vr1 interface lt-0/0/0.5
set routing-instances mk-vr1 routing-options static route 12.1.1.0/24 next-hop  10.0.1.2
set routing-instances mk-vr1 routing-options static route 14.1.1.0/24 next-hop  10.0.1.4
set routing-instances mk-vr1 routing-options static route 12.12.1.0/24 next-hop  10.0.1.1
set security zones security-zone ls-marketing-trust interfaces ge-0/0/6.1
set security zones security-zone ls-marketing-untrust interfaces lt-0/0/0.5
set security address-book marketing-external address design 12.1.1.0/24
set security address-book marketing-external address accounting 14.1.1.0/24
set security address-book marketing-external address others 12.12.1.0/24
set security address-book marketing-external address-set otherlsys address  design
set security address-book marketing-external address-set otherlsys address  accounting
set security address-book marketing-external attach zone ls-marketing-untrust
set security address-book marketing-internal address marketers 13.1.1.0/24
set security address-book marketing-internal attach zone ls-marketing-trust
set security policies from-zone ls-marketing-trust to-zone ls-marketing-untrust  policy permit-
all-to-otherlsys match source-address marketers
set security policies from-zone ls-marketing-trust to-zone ls-marketing-untrust  policy permit-
all-to-otherlsys match destination-address otherlsys
set security policies from-zone ls-marketing-trust to-zone ls-marketing-untrust  policy permit-
all-to-otherlsys match application any
set security policies from-zone ls-marketing-trust to-zone ls-marketing-untrust  policy permit-
all-to-otherlsys then permit
set security policies from-zone ls-marketing-untrust to-zone ls-marketing-trust policy permit-
all-from-otherlsys match source-address otherlsys
set security policies from-zone ls-marketing-untrust to-zone ls-marketing-trust  policy permit-
all-from-otherlsys match destination-address marketers
set security policies from-zone ls-marketing-untrust to-zone ls-marketing-trust  policy permit-
all-from-otherlsys match application any
```

```
set security policies from-zone ls-marketing-untrust to-zone ls-marketing-trust  policy permit-
all-from-otherlsys then permit
```

Step-by-Step Procedure

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

To configure a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

```
lsmarketingadmin1@host:ls-marketing-dept> configure
lsmarketingadmin1@host:ls-marketing-dept#
```

2. Configure the logical interface for a user logical system.

```
[edit interfaces]
lsmarketingadmin1@host:ls-marketing-dept# set ge-0/0/6 unit 1 family inet address 13.1.1.1/24
lsmarketingadmin1@host:ls-marketing-dept# set ge-0/0/6 unit 1 vlan-id 800
```

3. Configure the routing instance and assign interfaces.

```
[edit routing-instances]
lsmarketingadmin1@host:ls-marketing-dept# set mk-vr1 instance-type virtual-router
lsmarketingadmin1@host:ls-marketing-dept#  set mk-vr1 interface ge-0/0/6.1
lsmarketingadmin1@host:ls-marketing-dept#  set mk-vr1 interface lt-0/0/0.5
```

4. Configure static routes.

```
[edit routing-instances]
lsmarketingadmin1@host:ls-marketing-dept# set mk-vr1 routing-options static route
12.12.1.0/24 next-hop 10.0.1.2
lsmarketingadmin1@host:ls-marketing-dept# set mk-vr1 routing-options static route 14.1.1.0/24
next-hop 10.0.1.4
lsmarketingadmin1@host:ls-marketing-dept# set mk-vr1 routing-options static route
12.12.1.0/24 next-hop 10.0.1.1
```

5. Configure security zones and assign interfaces to each zone.

```
[edit security zones]
lsmarketingadmin1@host:ls-marketing-dept# set security-zone ls-marketing-trust interfaces
ge-0/0/6.1
lsmarketingadmin1@host:ls-marketing-dept# set security-zone ls-marketing-untrust interfaces
lt-0/0/0.5
```

6. Create address book entries.

```
[edit security]
lsmarketingadmin1@host:ls-marketing-dept# set address-book marketing-internal address
marketers 13.1.1.0/24
lsmarketingadmin1@host:ls-marketing-dept# set address-book marketing-external address design
12.1.1.0/24
lsmarketingadmin1@host:ls-marketing-dept# set address-book marketing-external address
accounting 14.1.1.0/24
lsmarketingadmin1@host:ls-marketing-dept# set address-book marketing-external address others
12.12.1.0/24
lsmarketingadmin1@host:ls-marketing-dept# set address-book marketing-external address-set
otherlsys address design
lsmarketingadmin1@host:ls-marketing-dept# set address-book marketing-external address-set
otherlsys address accounting
```

7. Attach address books to zones.

```
[edit security]
lsmarketingadmin1@host:ls-marketing-dept# set address-book marketing-internal attach zone ls-
marketing-trust
lsmarketingadmin1@host:ls-marketing-dept# set address-book marketing-external attach zone ls-
marketing-untrust
```

8. Configure a security policy that permits traffic from the ls-marketing-trust zone to the ls-marketing-untrust zone.

```
[edit security policies from-zone ls-marketing-trust to-zone ls-marketing-untrust]
lsmarketingadmin1@host:ls-marketing-dept# set policy permit-all-to-otherlsys match source-
address marketers
lsmarketingadmin1@host:ls-marketing-dept# set policy permit-all-to-otherlsys match
```

```
destination-address otherlsys
lsmarketingadmin1@host:ls-marketing-dept# set policy permit-all-to-otherlsys match
application any
lsmarketingadmin1@host:ls-marketing-dept# set policy permit-all-to-otherlsys then permit
```

9. Configure a security policy that permits traffic from the ls-marketing-untrust zone to the ls-marketing-trust zone.

```
[edit security policies from-zone ls-marketing-untrust to-zone ls-marketing-trust]
lsmarketingadmin1@host:ls-marketing-dept# set policy permit-all-from-otherlsys match source-
address otherlsys
lsmarketingadmin1@host:ls-marketing-dept# set policy permit-all-from-otherlsys match
destination-address marketers
lsmarketingadmin1@host:ls-marketing-dept# set policy permit-all-from-otherlsys match
application any
lsmarketingadmin1@host:ls-marketing-dept# set policy permit-all-from-otherlsys then permit
```

Results

From configuration mode, confirm your configuration by entering the `show routing-instances` and `show security` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
lsmarketingadmin1@host:ls-marketing-dept# show routing instances
    mk-vr1 {
        instance-type virtual-router;
        interface ge-0/0/6.1;
        interface lt-0/0/0.5;
        routing-options {
            static {
                route 12.1.1.0/24 next-hop 10.0.1.2;
                route 14.1.1.0/24 next-hop 10.0.1.4;
                route 12.12.1.0/24 next-hop 10.0.1.1;
            }
        }
    }
lsmarketingadmin1@host:ls-marketing-dept# show security
    address-book {
        marketing-external {
            address product-designers 12.1.1.0/24;
```

```
            address accounting 14.1.1.0/24;
            address others 12.12.1.0/24;
            address-set otherlsys {
                address product-designers;
                address accounting;
            }
            attach {
                zone ls-marketing-untrust;
            }
        }
        marketing-internal {
            address marketers 13.1.1.0/24;
            attach {
                zone ls-marketing-trust;
            }
        }
    }
        policies {
            from-zone ls-marketing-trust to-zone ls-marketing-untrust {
                policy permit-all-to-otherlsys {
                    match {
                        source-address marketers;
                        destination-address otherlsys;
                        application any;
                    }
                    then {
                        permit;
                    }
                }
            }
            from-zone ls-marketing-untrust to-zone ls-marketing-trust {
                policy permit-all-from-otherlsys {
                    match {
                        source-address otherlsys;
                        destination-address marketers;
                        application any;
                    }
                    then {
                        permit;
                    }
                }
            }
        }
```

```
        zones {
            security-zone ls-marketing-trust {
                interfaces {
                    ge-0/0/6.1;
                }
            }
            security-zone ls-marketing-untrust {
                interfaces {
                    lt-0/0/0.5;
                }
            }
        }
```

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

Configuring the ls-accounting-dept User Logical System

CLI Quick Configuration

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

```
set interfaces ge-0/0/7 unit 1 family inet address 14.1.1.1/24
set interfaces ge-0/0/7 unit 1 vlan-id 900
set routing-instances acct-vr1 instance-type virtual-router
set routing-instances acct-vr1 interface ge-0/0/7.1
set routing-instances acct-vr1 interface lt-0/0/0.7
set routing-instances acct-vr1 routing-options static route 12.12.1.0/24 next-hop 10.0.1.1
set routing-instances acct-vr1 routing-options static route 12.1.1.0/24 next-hop 10.0.1.2
set routing-instances acct-vr1 routing-options static route 13.1.1.0/24 next-hop 10.0.1.3
set security address-book accounting-internal address accounting 14.1.1.0/24
set security address-book accounting-internal attach zone ls-accounting-trust
set security address-book accounting-external address design 12.1.1.0/24
set security address-book accounting-external address marketing 13.1.1.0/24
set security address-book accounting-external address others 12.12.1.0/24
set security address-book accounting-external address-set otherlsys address design
set security address-book accounting-external address-set otherlsys address marketing
set security address-book accounting-external attach zone ls-accounting-untrust
set security policies from-zone ls-accounting-trust to-zone ls-accounting-untrust policy permit-
all-to-otherlsys match source-address accounting
set security policies from-zone ls-accounting-trust to-zone ls-accounting-untrust policy permit-
```

```
all-to-otherlsys match destination-address otherlsys
 set security policies from-zone ls-accounting-trust to-zone ls-accounting-untrust policy permit-
all-to-otherlsys match application any
set security policies from-zone ls-accounting-trust to-zone ls-accounting-untrust policy permit-
all-to-otherlsys then permit
set security policies from-zone ls-accounting-untrust to-zone ls-accounting-trust policy permit-
all-from-otherlsys match source-address otherlsys
set security policies from-zone ls-accounting-untrust to-zone ls-accounting-trust policy permit-
all-from-otherlsys match destination-address accounting
set security policies from-zone ls-accounting-untrust to-zone ls-accounting-trust policy permit-
all-from-otherlsys match application any
set security policies from-zone ls-accounting-untrust to-zone ls-accounting-trust policy permit-
all-from-otherlsys then permit
set security zones security-zone ls-accounting-trust interfaces ge-0/0/7.1
set security zones security-zone ls-accounting-untrust interfaces lt-0/0/0.7
```

Step-by-Step Procedure

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

To configure a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

```
lsaccountingadmin1@host:ls-accounting-dept> configure
lsaccountingadmin1@host:ls-accounting-dept#
```

2. Configure the logical interface for a user logical system.

```
[edit interfaces]
lsaccountingadmin1@host:ls-accounting-dept# set ge-0/0/7 unit 1 family inet address
14.1.1.1/24
lsaccountingadmin1@host:ls-accounting-dept# set ge-0/0/7 unit 1 vlan-id 900
```

3. Configure the routing instance and assign interfaces.

```
[edit routing-instances]
lsaccountingadmin1@host:ls-accounting-dept# set acct-vr1 instance-type virtual-router
```

```
lsaccountingadmin1@host:ls-accounting-dept#  set acct-vr1 interface ge-0/0/7.1
lsaccountingadmin1@host:ls-accounting-dept#  set acct-vr1 interface lt-0/0/0.7
```

4. Configure static routes.

```
[edit routing-instances]
lsaccountingadmin1@host:ls-accounting-dept# set acct-vr1 routing-options static route
12.1.1.0/24 next-hop 10.0.1.2
lsaccountingadmin1@host:ls-accounting-deptt# set acct-vr1 routing-options static route
13.1.1.0/24 next-hop 10.0.1.3
lsaccountingadmin1@host:ls-accounting-dept# set acct-vr1 routing-options static route
12.12.1.0/24 next-hop 10.0.1.1
```

5. Configure security zones and assign interfaces to each zone.

```
[edit security zones]
lsaccountingadmin1@host:ls-accounting-dept# set security-zone ls-accounting-trust interfaces
ge-0/0/7.1
lsaccountingadmin1@host:ls-accounting-dept# set security-zone ls-accounting-untrust
interfaces lt-0/0/0.7
```

6. Create address book entries.

```
[edit security]
lsaccountingadmin1@host:ls-accounting-dept# set address-book accounting-internal address
accounting 14.1.1.0/24
lsaccountingadmin1@host:ls-accounting-dept# set address-book accounting-external address
design 12.1.1.0/24
lsaccountingadmin1@host:ls-accounting-dept# set address-book accounting-external address
marketing 13.1.1.0/24
lsaccountingadmin1@host:ls-accounting-dept# set address-book accounting-external address
others 12.12.1.0/24
lsaccountingadmin1@host:ls-accounting-dept# set address-book accounting-external address-set
otherlsys address design
lsaccountingadmin1@host:ls-accounting-dept# set address-book accounting-external address-set
otherlsys address marketing
```

7. Attach address books to zones.

```
[edit security]
lsaccountingadmin1@host:ls-accounting-dept# set address-book accounting-internal attach zone
ls-accounting-trust
lsaccountingadmin1@host:ls-accounting-dept# set address-book accounting-external attach zone
ls-accounting-untrust
```

8. Configure a security policy that permits traffic from the ls-accounting-trust zone to the ls-accounting-untrust zone.

```
[edit security policies from-zone ls-accounting-trust to-zone ls-accounting-untrust]
lsaccountingadmin1@host:ls-accounting-dept# set policy permit-all-to-otherlsys match source-
address accounting
lsaccountingadmin1@host:ls-accounting-dept# set policy permit-all-to-otherlsys match
destination-address otherlsys
lsaccountingadmin1@host:ls-accounting-dept# set policy permit-all-to-otherlsys match
application any
lsaccountingadmin1@host:ls-accounting-dept# set policy permit-all-to-otherlsys then permit
```

9. Configure a security policy that permits traffic from the ls-accounting-untrust zone to the ls-accounting-trust zone.

```
[edit security policies from-zone ls-accounting-untrust to-zone ls-accounting-trust]
lsaccountingadmin1@host:ls-accounting-dept# set policy permit-all-from-otherlsys match source-
address otherlsys
lsaccountingadmin1@host:ls-accounting-dept# set policy permit-all-from-otherlsys match
destination-address accounting
lsaccountingadmin1@host:ls-accounting-dept# set policy permit-all-from-otherlsys match
application any
lsaccountingadmin1@host:ls-accounting-dept# set policy permit-all-from-otherlsys then permit
```

Results

From configuration mode, confirm your configuration by entering the `show routing-instances` and `show security` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
lsaccountingadmin1@host:ls-accounting-dept# show routing-instances
acct-vr1 {
    instance-type virtual-router;
    interface ge-0/0/7.1;
    interface lt-0/0/0.7;
    routing-options {
        static {
            route 12.12.1.0/24 next-hop 10.0.1.1;
            route 12.1.1.0/24 next-hop 10.0.1.2;
            route 13.1.1.0/24 next-hop 10.0.1.3;
        }
    }
}
lsaccountingadmin1@host:ls-accounting-dept# show security
address-book {
    accounting-internal {
        address accounting 14.1.1.0/24;
        attach {
            zone ls-accounting-trust;
        }
    }
    accounting-external {
        address design 12.1.1.0/24;
        address marketing 13.1.1.0/24;
        address others 12.12.1.0/24;
        address-set otherlsys {
            address design;
            address marketing;
        }
        attach {
            zone ls-accounting-untrust;
        }
    }
}
    policies {
        from-zone ls-accounting-trust to-zone ls-accounting-untrust {
```

```
            policy permit-all-to-otherlsys {
                match {
                    source-address accounting;
                    destination-address otherlsys;
                    application any;
                }
                then {
                    permit;
                }
            }
        }
        from-zone ls-accounting-untrust to-zone ls-accounting-trust {
            policy permit-all-from-otherlsys {
                match {
                    source-address otherlsys;
                    destination-address accounting;
                    application any;
                }
                then {
                    permit;
                }
            }
        }
    }
    zones {
        security-zone ls-accounting-trust {
            interfaces {
                ge-0/0/7.1;
            }
        }
        security-zone ls-accounting-untrust {
            interfaces {
                lt-0/0/0.7;
            }
        }
    }
```

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

Verification

IN THIS SECTION

- Verifying Policy Configuration | 167

To confirm that the configuration is working properly, perform these tasks:

Verifying Policy Configuration

Purpose

Verify information about policies and rules.

Action

From operational mode, enter the `show security policies detail` command to display a summary of all policies configured on the logical system.

SEE ALSO

User Logical Systems Configuration Overview | 51

Understanding Logical Systems Interfaces and Routing Instances | 113

Understanding Logical Systems Zones | 149

Understanding Logical Systems Security Policies | 213

Example: Configuring Security Zones for a User Logical Systems

IN THIS SECTION

- Requirements | 168
- Overview | 168

Configuration | 169

This example shows how to configure zones for a user logical system.

Requirements

Before you begin:

- Log in to the user logical system as the user logical system administrator. See "User Logical Systems Configuration Overview" on page 51.
- Use the `show system security-profile zones` command to see the zone resources allocated to the logical system.
- Logical interfaces for the user logical system must be configured. See "Example: Configuring Interfaces and Routing Instances for a User Logical Systems" on page 139.

Overview

This example configures the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System" on page 57.

This example creates the zones and address books described in Table 13 on page 168.

Table 13: User Logical System Zone and Address Book Configuration

Feature	Name	Configuration Parameters
Zones	ls-product-design-trust	• Bind to interface ge-0/0/5.1. • TCP reset enabled.
	ls-product-design-untrust	• Bind to interface lt-0/0/0.3.
Address books	product-design-internal	• Address product-designers: 12.1.1.0/24 • Attach to zone ls-product-design-trust

Table 13: User Logical System Zone and Address Book Configuration *(Continued)*

Feature	Name	Configuration Parameters
	product-design-external	• Address marketing: 13.1.1.0/24 • Address accounting: 14.1.1.0/24 • Address others: 12.12.1.0/24 • Address set otherlsys: marketing, accounting • Attach to zone ls-product-design-untrust

Configuration

IN THIS SECTION

- Procedure | **169**

Procedure

CLI Quick Configuration

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

```
set security address-book product-design-internal address product-designers  12.1.1.0/24
set security address-book product-design-internal attach zone  ls-product-design-trust
set security address-book product-design-external address marketing 13.1.1.0/24
set security address-book product-design-external address accounting 14.1.1.0/24
set security address-book product-design-external address others 12.12.1.0/24
set security address-book product-design-external address-set otherlsys address  marketing
set security address-book product-design-external address-set otherlsys address  accounting
set security address-book product-design-external attach zone  ls-product-design-untrust
set security zones security-zone ls-product-design-trust tcp-rst
```

```
set security zones security-zone ls-product-design-trust interfaces ge-0/0/5.1
set security zones security-zone ls-product-design-untrust interfaces lt-0/0/0.3
```

Step-by-Step Procedure

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

To configure zones in a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

```
lsdesignadmin1@host:ls-product-design> configure
lsdesignadmin1@host:ls-product-design#
```

2. Configure a security zone and assign it to an interface.

```
[edit security zones]
lsdesignadmin1@host:ls-product-design# set security-zone ls-product-design-trust interfaces
ge-0/0/5.1
```

3. Configure the TCP-Reset parameter for the zone.

```
[edit security zones security-zone ls-product-design-trust]
lsdesignadmin1@host:ls-product-design#  set tcp-rst
```

4. Configure a security zone and assign it to an interface.

```
[edit security zones]
lsdesignadmin1@host:ls-product-design# set security-zone ls-product-design-untrust interfaces
lt-0/0/0.3
```

5. Create global address book entries.

```
[edit security]
lsdesignadmin1@host:ls-product-design# set address-book product-design-internal address
product-designers 12.1.1.0/24
```

```
lsdesignadmin1@host:ls-product-design#  set address-book product-design-external address
marketing 13.1.1.0/24
lsdesignadmin1@host:ls-product-design#  set address-book product-design-external address
accounting 14.1.1.0/24
lsdesignadmin1@host:ls-product-design#  set address-book product-design-external address
others 12.12.1.0/24
lsdesignadmin1@host:ls-product-design#  set address-book product-design-external address-set
otherlsys address marketing
lsdesignadmin1@host:ls-product-design#  set address-book product-design-external address-set
otherlsys address accounting
```

6. Attach address books to zones.

```
[edit security]
lsdesignadmin1@host:ls-product-design# set address-book product-design-internal attach zone
ls-product-design-trust
lsdesignadmin1@host:ls-product-design#  set address-book product-design-external attach zone
ls-product-design-untrust
```

7.

Results

From configuration mode, confirm your configuration by entering the `show security` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
lsdesignadmin1@host:ls-product-design# show security
    address-book {
        product-design-internal {
            address product-designers 12.1.1.0/24;
            attach {
                zone ls-product-design-trust;
            }
        }
        product-design-external {
            address marketing 13.1.1.0/24;
            address accounting 14.1.1.0/24;
            address others 12.12.1.0/24;
            address-set otherlsys {
                address marketing;
```

```
                address accounting;
            }
            attach {
                zone ls-product-design-untrust;
            }
        }
    }
        zones {
            security-zone ls-product-design-trust {
                tcp-rst;
                interfaces {
                    ge-0/0/5.1;
                }
            }
            security-zone ls-product-design-untrust {
                interfaces {
                    lt-0/0/0.3;
                }
            }
        }
```

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

SEE ALSO

- Understanding Logical Systems Zones | 149
- User Logical Systems Configuration Overview | 51

RELATED DOCUMENTATION

- Example: Configuring Security Policies in a User Logical Systems | 216

User Authentication for Logical Systems

IN THIS SECTION

- Example: Configuring Access Profiles (Primary Administrators Only) | 173
- Example: Configuring Security Features for the Primary Logical Systems | 176
- Understanding Logical System Firewall Authentication | 185
- Example: Configuring Firewall Authentication for a User Logical System | 187
- Understanding Integrated User Firewall support in a Logical System | 193
- Example: Configuring Integrated User Firewall Identification Management for a User Logical System | 194
- Example: Configure Integrated User Firewall in Customized Model for Logical System | 205

User authentication for logical systems enables you to define firewall users and create policies that require the users to authenticate themselves through one of two authentication schemes: pass-through authentication or web authentication. For more information, see the following topics:

Example: Configuring Access Profiles (Primary Administrators Only)

IN THIS SECTION

- Requirements | 173
- Overview | 174
- Configuration | 174

The primary administrator is responsible for configuring access profiles in the primary logical system. This example shows how to configure access profiles.

Requirements

Before you begin:

- Log in to the primary logical system as the primary administrator. See "Understanding the Primary Logical Systems and the Primary Administrator Role" on page 23.
- Read *Firewall User Authentication Overview*.

Overview

This example configures an access profile for LDAP authentication for logical system users. This example creates the access profile described in Table 14 on page 174.

NOTE: The primary administrator creates the access profile.

Table 14: Access Profile Configuration

Name	Configuration Parameters

ldap1	• LDAP is used as the first (and only) authentication method.
	• Base distinguished name:
	  • Organizational unit name (OU): people
	  • Domain components (DC): example, com
	• A user's LDAP distinguished name is assembled through the use of a common name identifier, username, and base distinguished name. The common name identifier is user ID (UID).
	• The LDAP server address is 10.155.26.104 and is reached through port 389.

Configuration

IN THIS SECTION

Procedure | 175

Procedure

CLI Quick Configuration

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

NOTE: You must be logged in as the primary administrator.

```
set access profile ldap1 authentication-order ldap
set access profile ldap1 ldap-options base-distinguished-name ou=people,dc=example,dc=com
set access profile ldap1 ldap-options assemble common-name uid
set access profile ldap1 ldap-server 10.155.26.104 port 389
```

Step-by-Step Procedure

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

To configure an access profile in the primary logical system:

1. Log in to the primary logical system as the primary administrator and enter configuration mode.

```
admin@host> configure
admin@host#
```

2. Configure an access profile and set the authentication order.

```
[edit access profile ldap1]
admin@host# set authentication-order ldap
```

3. Configure LDAP options.

```
[edit access profile ldap1]
admin@host# set ldap-options base-distinguished-name ou=people,dc=example,dc=com
admin@host# set ldap-options assemble common-name uid
```

4. Configure the LDAP server.

```
[edit access profile ldap1]
admin@host# set ldap-server 10.155.26.104 port 389
```

Results

From configuration mode, confirm your configuration by entering the `show access profile` *profile-name* command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
admin@host# show access profile ldap1
authentication-order ldap;
ldap-options {
    base-distinguished-name ou=people,dc=example,dc=com;
    assemble {
        common-name uid;
    }
}
    ldap-server {
        10.155.26.104 port 389;
    }
```

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

SEE ALSO

User Logical Systems Configuration Overview | 51

Example: Configuring Security Features for the Primary Logical Systems

IN THIS SECTION

- Requirements | 177
- Overview | 177

Configuration | 179
Verification | 184

This example shows how to configure security features, such as zones, policies, and firewall authentication, for the primary logical system.

Requirements

Before you begin:

- Log in to the primary logical system as the primary administrator. See "Example: Configuring Root Password for Logical Systems" on page 55.
- Use the `show system security-profile` command to see the resources allocated to the primary logical system.
- Configure logical interfaces for the primary logical system. See "Example: Configuring Interfaces, Routing Instances, and Static Routes for the Primary and Interconnect Logical Systems and Logical Tunnel Interfaces for the User Logical Systems (Primary Administrators Only)" on page 114.
- Configure the access profile ldap1 in the primary logical system. The ldap1 access profile is used for Web authentication of firewall users.

Overview

In this example, you configure security features for the primary logical system, called root-logical-system, shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System" on page 57. This example configures the security features described in Table 15 on page 177.

Table 15: root-logical-system Security Feature Configuration

Feature	Name	Configuration Parameter
Zones	ls-root-trust	Bind to interface ge-0/0/4.0.
	ls-root-untrust	Bind to interface lt-0/0/0.1

Table 15: root-logical-system Security Feature Configuration *(Continued)*

Feature	Name	Configuration Parameter
Address books	root-internal	• Address primaries: 10.12.12.0/24 • Attach to zone ls-root-trust
	root-external	• Address design: 10.12.1.0/24 • Address accounting: 10.14.1.0/24 • Address marketing: 10.13.1.0/24 • Address set userlsys: design, accounting, marketing • Attach to zone ls-root-untrust
Security policies	permit-to-userlsys	Permit the following traffic: • From zone: ls-root-trust • To zone: ls-root-untrust • Source address: primaries • Destination address: userlsys • Application: any
	permit-authorized-users	Permit the following traffic: • From zone: ls-root-untrust • To zone: ls-root-trust • Source address: userlsys • Destination address: primaries • Application: junos-http, junos-https

Table 15: root-logical-system Security Feature Configuration *(Continued)*

Feature	Name	Configuration Parameter
Firewall authentication		• Web authentication • Authentication success banner “WEB AUTH LOGIN SUCCESS” • Default access profile ldap1
HTTP daemon		Activate on interface ge-0/0/4.0

Configuration

IN THIS SECTION

- Procedure | 179

Procedure

CLI Quick Configuration

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

```
set security address-book root-internal address masters 10.12.12.0/24
set security address-book root-internal attach zone ls-root-trust
set security address-book root-external address design 10.12.1.0/24
set security address-book root-external address accounting 10.14.1.0/24
set security address-book root-external address marketing 10.13.1.0/24
set security address-book root-external address-set userlsys address design
set security address-book root-external address-set userlsys address accounting
set security address-book root-external address-set userlsys address marketing
set security address-book root-external attach zone ls-root-untrust
set security policies from-zone ls-root-trust to-zone ls-root-untrust policy permit-to-userlsys
```

```
match source-address masters
set security policies from-zone ls-root-trust to-zone ls-root-untrust policy permit-to-userlsys
match destination-address userlsys
set security policies from-zone ls-root-trust to-zone ls-root-untrust policy permit-to-userlsys
match application any
set security policies from-zone ls-root-trust to-zone ls-root-untrust policy permit-to-userlsys
then permit
set security policies from-zone ls-root-untrust to-zone ls-root-trust policy permit-authorized-
users match source-address userlsys
set security policies from-zone ls-root-untrust to-zone ls-root-trust policy permit-authorized-
users match destination-address masters
set security policies from-zone ls-root-untrust to-zone ls-root-trust policy permit-authorized-
users match application junos-http
set security policies from-zone ls-root-untrust to-zone ls-root-trust policy permit-authorized-
users match application junos-https
set security policies from-zone ls-root-untrust to-zone ls-root-trust policy permit-authorized-
users then permit firewall-authentication web-authentication
set security zones security-zone ls-root-trust interfaces ge-0/0/4.0
set security zones security-zone ls-root-untrust interfaces lt-0/0/0.1
set system services web-management http interface ge-0/0/4.0
set access firewall-authentication web-authentication default-profile ldap1
set access firewall-authentication web-authentication banner success "WEB AUTH LOGIN SUCCESS"
```

Step-by-Step Procedure

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

To configure zones and policies for the primary logical system:

1. Log in to the primary logical system as the primary administrator and enter configuration mode.

```
admin@host> configure
admin@host#
```

2. Create security zones and assign interfaces to each zone.

```
[edit security zones]
admin@host# set security-zone ls-root-trust interfaces ge-0/0/4.0
admin@host# set security-zone ls-root-untrust interfaces lt-0/0/0.1
```

3. Create address book entries.

```
[edit security]
admin@host# set address-book root-internal address masters 10.12.12.0/24
admin@host# set address-book root-external address design 10.12.1.0/24
admin@host# set address-book root-external address accounting 10.14.1.0/24
admin@host# set address-book root-external address marketing 10.13.1.0/24
admin@host# set address-book root-external address-set userlsys address design
admin@host# set address-book root-external address-set userlsys address accounting
admin@host# set address-book root-external address-set userlsys address marketing
```

4. Attach address books to zones.

```
[edit security]
admin@host# set address-book root-internal attach zone ls-root-trust
admin@host# set address-book root-external attach zone ls-root-untrust
```

5. Configure a security policy that permits traffic from the ls-root-trust zone to the ls-root-untrust zone.

```
[edit security policies from-zone ls-root-trust to-zone ls-root-untrust]
admin@host# set policy permit-to-userlsys match source-address masters
admin@host# set policy permit-to-userlsys match destination-address userlsys
admin@host# set policy permit-to-userlsys match application any
admin@host# set policy permit-to-userlsys then permit
```

6. Configure a security policy that authenticates traffic from the ls-root-untrust zone to the ls-root-trust zone.

```
[edit security policies from-zone ls-root-untrust to-zone ls-root-trust]
admin@host# set policy permit-authorized-users match source-address userlsys
admin@host# set policy permit-authorized-users match destination-address masters
admin@host# set policy permit-authorized-users match application junos-http
```

```
admin@host# set policy permit-authorized-users match application junos-https
admin@host# set policy permit-authorized-users then permit firewall-authentication web-
authentication
```

7. Configure the Web authentication access profile and define a success banner.

```
[edit access]
admin@host# set firewall-authentication web-authentication default-profile ldap1
admin@host# set firewall-authentication web-authentication banner success “WEB AUTH LOGIN
SUCCESS”
```

8. Activate the HTTP daemon on the device.

```
[edit system]
admin@host# set services web-management http interface ge-0/0/4.0
```

Results

From configuration mode, confirm your configuration by entering the `show security`, `show access`, and `show system services` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this `show` command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
admin@host# show security
...
address-book {
    root-internal {
        address masters 10.12.12.0/24;
        attach {
            zone ls-root-trust;
        }
    }
    root-external {
        address design 10.12.1.0/24;
        address accounting 10.14.1.0/24;
        address marketing 10.13.1.0/24;
        address-set userlsys {
```

```
            address design;
            address accounting;
            address marketing;
        }
        attach {
            zone ls-root-untrust;
        }
    }
}
    policies {
        from-zone ls-root-trust to-zone ls-root-untrust {
            policy permit-to-userlsys {
                match {
                    source-address masters;
                    destination-address userlsys;
                    application any;
                }
                then {
                    permit;
                }
            }
        }
        from-zone ls-root-untrust to-zone ls-root-trust {
            policy permit-authorized-users {
                match {
                    source-address userlsys;
                    destination-address masters;
                    application [ junos-http junos-https ];
                }
                then {
                    permit {
                        firewall-authentication {
                            web-authentication;
                        }
                    }
                }
            }
        }
    }
    zones {
        security-zone ls-root-trust {
            interfaces {
                ge-0/0/4.0;
```

```
            }
        }
        security-zone ls-root-untrust {
            interfaces {
                lt-0/0/0.1;
            }
        }
    }
[edit]
admin@host# show access
...
firewall-authentication {
    web-authentication {
        default-profile ldap1;
        banner {
            success "WEB AUTH LOGIN SUCCESS";
        }
    }
}
[edit]
admin@host# show system services
web-management {
    http {
        interface ge-0/0/4.0;
    }
}
```

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

Verification

IN THIS SECTION

- Verifying Policy Configuration | 185

To confirm that the configuration is working properly, perform these tasks:

Verifying Policy Configuration

Purpose

Verify information about policies and rules.

Action

From operational mode, enter the `show security policies detail` command to display a summary of all policies configured on the logical system.

SEE ALSO

Understanding Logical Systems Zones | 149

Understanding Logical Systems Security Policies | 213

Understanding Logical System Firewall Authentication

A firewall user is a network user who must provide a username and password for authentication when initiating a connection across the firewall. Junos OS enables administrators to restrict and permit firewall users to access protected resources (different zones) behind a firewall based on their source IP address and other credentials.

The primary administrator is responsible for configuring access profiles in the primary logical system. Access profiles store usernames and passwords of users or point to external authentication servers where such information is stored. Access profiles configured at the primary logical system are available to all user logical systems.

The primary administrator configures the maximum and reserved numbers of firewall authentications for each user logical system. The user logical system administrator can then create firewall authentications in the user logical system. From a user logical system, the user logical system administrator can use the `show system security-profile auth-entry` command to view the number of authentication resources allocated to the user logical system.

To configure the access profile, the primary administrator uses the `profile` *configuration statement* at the [`edit access`] hierarchy level in the primary logical system. The access profile can also include the order of authentication methods, LDAP or RADIUS server options, and session options.

The user logical system administrator can then associate the access profile with a security policy in the user logical system. The user logical system administrator also specifies the type of authentication:

- With pass-through authentication, a host or a user from one zone tries to access resources on another zone using an FTP, a telnet, or an HTTP client. The device uses FTP, Telnet, or HTTP to collect username and password information, and subsequent traffic from the user or host is allowed or denied based on the result of this authentication.
- With Web authentication, users use HTTP to connect to an IP address on the device that is enabled for Web authentication and are prompted for the username and password. Subsequent traffic from the user or host to the protected resource is allowed or denied based on the result of this authentication.

The user logical system administrator configures the following properties for firewall authentication in the user logical system:

- Security policy that specifies firewall authentication for matching traffic. Firewall authentication is specified with the `firewall-authentication` configuration statement at the [`edit security policies from-zone` *`zone-name`* `to-zone` *`zone-name`* `policy` *`policy-name`* `then permit`] hierarchy level.

  Users or user groups in an access profile who are allowed access by the policy can optionally be specified with the client-match configuration statement. (If no users or user groups are specified, any user who is successfully authenticated is allowed access.)

  For pass-through authentication, the access profile can optionally be specified and Web redirect (redirecting the client system to a webpage for authentication) can be enabled.

- Type of authentication (pass-through or Web authentication), default access profile, and success banner for the FTP, Telnet, or HTTP session. These properties are configured with the `firewall-authentication` configuration statement at the [`edit access`] hierarchy level.
- Host inbound traffic. Protocols, services, or both are allowed to access the logical system. The types of traffic are configured with the `host-inbound-traffic` configuration statement at the [`edit security zones security-zone` *`zone-name`*] or [`edit security zones security-zone` *`zone-name`* `interfaces` *`interface-name`*] hierarchy levels.

From a user logical system, the user logical system administrator can use the `show security firewall-authentication users` or `show security firewall-authentication history` commands to view the information about firewall users and history for the user logical system. From the primary logical system, the primary administrator can use the same commands to view information for the primary logical system, a specific user logical system, or all logical systems.

SEE ALSO

User Logical Systems Configuration Overview | 51

Understanding Logical Systems Security Profiles (Primary Administrators Only) | 71

Firewall User Authentication Overview

Example: Configuring Firewall Authentication for a User Logical System

IN THIS SECTION

- Requirements | 187
- Overview | 187
- Configuration | 188
- Verification | 192

This example shows how to configure firewall authentication for a user logical system.

Requirements

Before you begin:

- Log in to the user logical system as the logical system administrator. See "User Logical Systems Configuration Overview" on page 51.
- Use the `show system security-profiles auth-entry` command to see the firewall authentication entries allocated to the logical system.
- Access profiles must be configured in the primary logical system by the primary administrator.

Overview

This example configures the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System" on page 57.

In this example, users in the ls-marketing-dept and ls-accounting-dept logical systems are required to authenticate when initiating certain connections to the product designers subnet. This example configures the firewall authentication described in Table 16 on page 188.

NOTE: This example uses the access profile configured and address book entries configured in "Example: Configuring Security Zones for a User Logical Systems" on page 167.

Table 16: User Logical System Firewall Authentication Configuration

Feature	Name	Configuration Parameters
Security policy	permit-authorized-users **NOTE**: Policy lookup is performed in the order that the policies are configured. The first policy that matches the traffic is used. If you have previously configured a policy that permits traffic for the same from-zone, to-zone, source address, and destination address but with application any, the policy configured in this example would never be matched. (See "Example: Configuring Security Policies in a User Logical Systems" on page 216.) Therefore, this policy should be reordered so that it is checked first.	Permit firewall authentication for the following traffic: • From zone: ls-product-design-untrust • To zone: ls-product-design-trust • Source address: otherlsys • Destination address: product-engineers • Application: junos-h323 The ldap1 access profile is used for pass-through authentication.
Firewall authentication		• Pass-through authentication • HTTP login prompt “welcome” • Default access profile ldap1

Configuration

IN THIS SECTION

• Procedure | 189

Procedure

CLI Quick Configuration

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

```
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust policy
permit-authorized-users match source-address otherlsys
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust policy
permit-authorized-users match destination-address product-designers
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust policy
permit-authorized-users match application junos-h323
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust policy
permit-authorized-users then permit firewall-authentication pass-through access-profile ldap1
set access firewall-authentication pass-through default-profile ldap1
set access firewall-authentication pass-through http banner login “welcome”
```

Step-by-Step Procedure

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

To configure firewall authentication in a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

```
lsdesignadmin1@host:ls-product-design> configure
lsdesignadmin1@host:ls-product-design#
```

2. Configure a security policy that permits firewall authentication.

```
[edit security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust]
lsdesignadmin1@host:ls-product-design# set policy permit-authorized-users match source-
address otherlsys
lsdesignadmin1@host:ls-product-design# set policy permit-authorized-users match destination -
address product-designers
lsdesignadmin1@host:ls-product-design# set policy permit-authorized-users match application
```

```
junos-h323
lsdesignadmin1@host:ls-product-design# set policy permit-authorized-users then permit
firewall-authentication pass-through access-profile ldap1
```

3. Reorder the security policies.

```
[edit]
lsdesignadmin1@host:ls-product-design# insert security policies from-zone ls-product-design-
untrust to-zone ls-product-design-trust policy permit-authorized-users before policy permit-
all-from-otherlsys
```

4. Configure firewall authentication.

```
[edit access firewall-authentication]
lsdesignadmin1@host:ls-product-design#  set pass-through http banner login "welcome"
lsdesignadmin1@host:ls-product-design#  set pass-through default-profile ldap1
```

Results

From configuration mode, confirm your configuration by entering the `show security policies` and `show access firewall-authentication` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
lsdesignadmin1@host:ls-product-design# show security policies
from-zone ls-product-design-trust to-zone ls-product-design-untrust {
    policy permit-all-to-otherlsys {
        match {
            source-address product-designers;
            destination-address otherlsys;
            application any;
        }
        then {
            permit;
        }
    }
}
    from-zone ls-product-design-untrust to-zone ls-product-design-trust {
        policy permit-authorized-users {
            match {
```

```
                source-address otherlsys;
                destination-address product-designers;
                application junos-h323;
            }
            then {
                permit {
                    firewall-authentication {
                        pass-through {
                            access-profile ldap1;
                        }
                    }
                }
            }
        }
        policy permit-all-from-otherlsys {
            match {
                source-address otherlsys;
                destination-address product-designers;
                application any;
            }
            then {
                permit;
            }
        }
    }
lsdesignadmin1@host:ls-product-design# show access firewall-authentication
pass-through {
    default-profile ldap1;
    http {
        banner {
            login welcome;
        }
    }
}
```

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

Verification

IN THIS SECTION

- Verifying Firewall User Authentication and Monitoring Users and IP Addresses | 192

To confirm that the configuration is working properly, perform these tasks:

Verifying Firewall User Authentication and Monitoring Users and IP Addresses

Purpose

Display firewall authentication user history and verify the number of firewall users who successfully authenticated and firewall users who failed to log in.

Action

From operational mode, enter these `show` commands.

```
lsdesignadmin1@host:ls-product-design> show security firewall-authentication history
lsdesignadmin1@host:ls-product-design> show security firewall-authentication history identifier
id
lsdesignadmin1@host:ls-product-design> show security firewall-authentication users
lsdesignadmin1@host:ls-product-design> show security firewall-authentication users identifier id
```

SEE ALSO

User Logical Systems Configuration Overview | 51

Example: Configuring Pass-Through Authentication

Understanding Integrated User Firewall support in a Logical System

IN THIS SECTION

- Limitation of Using User Firewall Authentication | 194
- Limitation of Using User Firewall Authentication in Customized Model on Logical Systems | 194

Starting in Junos OS Release 18.3R1, the support for authentication sources is extended to include Local authentication, Active Directory (AD) authentication, and firewall authentication in addition to the existing support for authentication sources Juniper Identity Management Service (JIMS) and ClearPass authentication.

Starting in Junos OS Release 18.2R1, the support for user firewall authentication is enhanced using a shared model. In this model, user logical systems share user firewall configuration and authentication entries with the primary logical system and the integrated user firewall authentication is supported in a user logical system.

In the shared model, user firewall related configuration is configured under the primary logical system, such as authentication source, authentication source priority, authentication entries timeout, and IP query or Individual query and so on. The user firewall provides user information service for an application in the SRX Series device, such as policy and logging. Traffic from a user logical system queries authentication tables from the primary logical system.

The authentication tables are managed by a primary logical system. The user logical systems share the authentication tables. Traffic from the primary logical system and the user logical systems query the same authentication table. User logical systems enable the use of the source-identity in security policy.

For example, if the primary logical system is configured with **employee** and the user logical system is configured with the source-identity **manager**, then the reference group of this authentication entry includes **employee** and **manager**. This reference group contains the same authentication entries from primary logical system and user logical system.

Starting in Junos OS Release 19.3R1, support for user firewall authentication is enhanced by using a customized model through integrated JIMS with active mode. In this model, the logical system extracts the authentication entries from the root level. The primary logical system is configured to the JIMS server based on the logical system and tenant system name. In active mode the SRX series device actively queries the authenticaton entries received from the JIMS server through HTTPs protocol. To reduce the data exchange, firewall filters are applied.

The user firewall uses the logical system name as a differentiator and is consistent between the JIMS server and SRX series device. The JIMS server sends the differentiator which is included in the

authentication entry. The authentication entries are distributed into the root logical system, when the differentiator is set as default for primary logical system.

The user firewall supports In-service software upgrade (ISSU) for logical systems, as user firewall changes the internal database table format from Junos OS Release 19.2R1 onwards. Prior to Junos OS Release 19.2R1, ISSU is not supported for logical systems.

Limitation of Using User Firewall Authentication

Using user firewall authentication on tenant systems has the following limitation:

- The authentication entries are collected by the JIMS server based on the IP address from the customer network. If the IP addresses overlap, then the authentication entry changes when users log in under different user logical systems.

Limitation of Using User Firewall Authentication in Customized Model on Logical Systems

Using user firewall authentication in customized model on logical systems has the following limitation:

- The JIMS server configurations to be configured under the root logical systems.
- The logical system name should be consistent and unique between the JIMS server and the SRX series device.

SEE ALSO

show services user-identification authentication-table

Example: Configuring Integrated User Firewall Identification Management for a User Logical System

IN THIS SECTION

- Requirements | **195**
- Overview | **196**
- Configuration | **196**

Verification | 202

This example shows how to configure the SRX Series device's advanced query feature for obtaining user identity information from the Juniper Identity Management Service (JIMS) and the security policy to match the source identity for a user logical system. In the root logical system, user firewall is configured with JIMS, and then the root logical system manages all of authentication entries coming from JIMS. In this example, all of user logical systems share their authentication entries with the root logical system.

Requirements

This example uses the following hardware and software components:

- SRX1500 devices operating in chassis clustering
- JIMS server
- Junos OS Release 18.2 R1

Before you begin:

- Log in to the user logical system as the logical system administrator. See "User Logical Systems Overview" on page 51
- Configure user logical systems lsys1 and lsys2. See Example: Configuring User Logical Systems
- Configure security profile on primary logical system and assign it to user logical systems lsys1 and lsys2. See Example: Configuring Logical Systems Security Profiles (Primary Administrators Only)
- Configure interfaces and routing options on logical systems root logical system, user logical systems lsys1, and lsys2. See Example: Configuring Interfaces, Routing Instances, and Static Routes for the Primary and Interconnect Logical Systems and Logical Tunnel Interfaces for the User Logical Systems (Primary Administrators Only) and Example: Configuring Interfaces and Routing Instances for a User Logical Systems
- Configure security policies for a user logical systems. See Example: Configuring Security Policies in a User Logical Systems
- Configure zones for a user logical system. See Example: Configuring Security Zones for a User Logical Systems
- Configure logical systems in a basic active/passive chassis cluster. See Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (Primary Administrators Only)

Overview

In this example, you can configure JIMS with HTTPs connection on port 443 and primary server with IPv4 address on primary logical system, policy p1 with source-identity "group1" of dc0 domain on logical system lsys1, policy p1 with source-identity "group1" of dc0 domain on logical system lsys2, and send traffic from and through logical system lsys1 to logical system lsys2. You can view the authentication entries on primary logical system and user logical systems (lsys1 and lsys2) even after rebooting the primary node.

Configuration

IN THIS SECTION

- CLI Quick Configuration | **196**
- Configuring user firewall identification management | **198**
- Results | **201**

CLI Quick Configuration

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

```
set logical-systems lsys1 security policies from-zone lsys1_trust to-zone lsys1_trust policy
lsys1_policy1 match source-address any
set logical-systems lsys1 security policies from-zone lsys1_trust to-zone lsys1_trust policy
lsys1_policy1 match destination-address any
set logical-systems lsys1 security policies from-zone lsys1_trust to-zone lsys1_trust policy
lsys1_policy1 match application any
set logical-systems lsys1 security policies from-zone lsys1_trust to-zone lsys1_trust policy
lsys1_policy1 then permit
set logical-systems lsys1 security policies from-zone lsys1_trust to-zone lsys1_trust policy
lsys1_policy1 match source-identity "example.com\group1"
set logical-systems lsys1 security policies from-zone lsys1_trust to-zone lsys1_trust policy
lsys1_policy1 then permit
set logical-systems lsys1 security policies from-zone lsys1_trust to-zone lsys1_untrust policy
lsys1_policy2 match source-address any
set logical-systems lsys1 security policies from-zone lsys1_trust to-zone lsys1_untrust policy
lsys1_policy2 match destination-address any
```

```
set logical-systems lsys1 security policies from-zone lsys1_trust to-zone lsys1_untrust policy
lsys1_policy2 match application any
set logical-systems lsys1 security policies from-zone lsys1_trust to-zone lsys1_untrust policy
lsys1_policy2 then permit
set logical-systems lsys1 security policies from-zone lsys1_untrust to-zone lsys1_trust policy
lsys1_policy3 match source-address any
set logical-systems lsys1 security policies from-zone lsys1_untrust to-zone lsys1_trust policy
lsys1_policy3 match destination-address any
set logical-systems lsys1 security policies from-zone lsys1_untrust to-zone lsys1_trust policy
lsys1_policy3 match application any
set logical-systems lsys1 security policies from-zone lsys1_untrust to-zone lsys1_trust policy
lsys1_policy3 then permit
set logical-systems lsys1 security policies policy-rematch
set logical-systems lsys2 security policies from-zone lsys2_untrust to-zone lsys2_untrust policy
lsys2_policy1 match source-address any
set logical-systems lsys2 security policies from-zone lsys2_untrust to-zone lsys2_untrust policy
lsys2_policy1 match destination-address any
set logical-systems lsys2 security policies from-zone lsys2_untrust to-zone lsys2_untrust policy
lsys2_policy1 match application any
set logical-systems lsys2 security policies from-zone lsys2_untrust to-zone lsys2_untrust policy
lsys2_policy1 match source-identity "example.com\group2"
set logical-systems lsys2 security policies from-zone lsys2_untrust to-zone lsys2_untrust policy
lsys2_policy1 then permit
set logical-systems lsys2 security policies policy-rematch
set services user-identification identity-management connection connect-method https
set services user-identification identity-management connection port 443
set services user-identification identity-management connection primary address 192.0.2.5
set services user-identification identity-management connection primary client-id otest
set services user-identification identity-management connection primary client-secret "$ABC123"
set security policies from-zone root_trust to-zone root_trust policy root_policy1 match source-
address any
set security policies from-zone root_trust to-zone root_trust policy root_policy1 match
destination-address any
set security policies from-zone root_trust to-zone root_trust policy root_policy1 match
application any
set security policies from-zone root_trust to-zone root_trust policy root_policy1 then permit
set security policies policy-rematch
set security zones security-zone root_trust interfaces reth1.0 host-inbound-traffic system-
services all
set security zones security-zone root_trust interfaces reth1.0 host-inbound-traffic protocols all
set security zones security-zone root_trust interfaces lt-0/0/0.1 host-inbound-traffic system-
services all
set security zones security-zone root_trust interfaces lt-0/0/0.1 host-inbound-traffic protocols
```

```
all
set firewall family inet filter impair-ldap term allow_all then accept
```

Configuring user firewall identification management

Step-by-Step Procedure

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

To configure user firewall identification management:

1. Log in to the primary logical system as the primary administrator and enter configuration mode.

```
user@host> configure
user@host#
```

2. Create logical systems.

```
[edit logical-systems]
user@host#set LSYS0
user@host#set LSYS1
user@host#set LSYS2
```

3. Configure a security policy lsys1_policy1 with source-identity group1 on logical system lsys1 that permits traffic from lsys1_trust to lsys1_trust.

```
[edit security policies]
user@host#set from-zone lsys1_trust to-zone lsys1_trust policy lsys1_policy1 match source-
address any
user@host#set from-zone lsys1_trust to-zone lsys1_trust policy lsys1_policy1 match
destination-address any
user@host#set from-zone lsys1_trust to-zone lsys1_trust policy lsys1_policy1 match
application any
user@host#set from-zone lsys1_trust to-zone lsys1_trust policy lsys1_policy1 match source-
identity "example.com\group1"
user@host#set from-zone lsys1_trust to-zone lsys1_trust policy lsys1_policy1 then permit
```

4. Configure a security policy lsys1_policy2 that permits traffic from lsys1_trust to lsys1_untrust.

```
[edit security policies]
user@host#set from-zone lsys1_trust to-zone lsys1_untrust policy lsys1_policy2 match source-address any
user@host#set from-zone lsys1_trust to-zone lsys1_untrust policy lsys1_policy2 match destination-address any
user@host#set from-zone lsys1_trust to-zone lsys1_untrust policy lsys1_policy2 match application any
user@host#set from-zone lsys1_trust to-zone lsys1_untrust policy lsys1_policy2 then permit
```

5. Configure a security policy lsys1_policy3 that permits traffic from lsys1_untrust to lsys1_trust.

```
[edit security policies]
user@host#set from-zone lsys1_untrust to-zone lsys1_trust policy lsys1_policy3 match source-address any
user@host#set from-zone lsys1_untrust to-zone lsys1_trust policy lsys1_policy3 match destination-address any
user@host#set from-zone lsys1_untrust to-zone lsys1_trust policy lsys1_policy3 match application any
user@host#set from-zone lsys1_untrust to-zone lsys1_trust policy lsys1_policy3 then permit
user@host#set policy-rematch
```

6. Configure security zone and assign interfaces to each zone.

```
[edit security zones]
user@host#set security-zone lsys1_trust interfaces reth2.0 host-inbound-traffic system-services all
user@host#set security-zone lsys1_trust interfaces reth2.0 host-inbound-traffic protocols all
user@host#set security-zone lsys1_trust interfaces lt-0/0/0.11 host-inbound-traffic system-services all
user@host#set security-zone lsys1_trust interfaces lt-0/0/0.11 host-inbound-traffic protocols all
user@host#set security-zone lsys1_untrust interfaces reth3.0 host-inbound-traffic system-services all
user@host#set security-zone lsys1_untrust interfaces reth3.0 host-inbound-traffic protocols all
```

7. Configure a security policy lsys2_policy1 with source-identity group1 that permits traffic from lsys2_untrust to lsys2_untrust on lsys2.

```
[edit security policies]
user@host#set from-zone lsys2_untrust to-zone lsys2_untrust policy lsys2_policy1 match
source-address any
user@host#set from-zone lsys2_untrust to-zone lsys2_untrust policy lsys2_policy1 match
destination-address any
user@host#set from-zone lsys2_untrust to-zone lsys2_untrust policy lsys2_policy1 match
application any
user@host#set from-zone lsys2_untrust to-zone lsys2_untrust policy lsys2_policy1 match
source-identity "example.com\group2"
user@host#set from-zone lsys2_untrust to-zone lsys2_untrust policy lsys2_policy1 then permit
user@host#set policy-rematch
```

8. Configure security zones and assign interfaces to each zone on lsys2.

```
[edit security zones]
user@host#set security-zone lsys2_untrust interfaces reth4.0 host-inbound-traffic system-
services all
user@host#set security-zone lsys2_untrust interfaces reth4.0 host-inbound-traffic protocols
all
user@host#set security-zone lsys2_untrust interfaces lt-0/0/0.21 host-inbound-traffic
system-services all
user@host#set security-zone lsys2_untrust interfaces lt-0/0/0.21 host-inbound-traffic
protocols all
```

9. Configure JIMS as the authentication source for advanced query requests with the primary address. The SRX Series device requires this information to contact the server.

```
[edit services user-identification identity-management]
user@host#set connection port 443
user@host#set connection connect-method https
user@host#set connection primary address 192.0.2.5
user@host#set connection primary client-id otest
user@host#set connection primary client-secret test
user@host#set authentication-entry-timeout 0
```

10. Configure security policies and zones on primary logical system.

```
[edit security policies]
user@host#set from-zone root_trust to-zone root_trust policy root_policy1 match source-
address any
user@host#set from-zone root_trust to-zone root_trust policy root_policy1 match destination-
address any
user@host#set from-zone root_trust to-zone root_trust policy root_policy1 match application
any
user@host#set from-zone root_trust to-zone root_trust policy root_policy1 then permit
user@host#set policy-rematch
```

11. Configure security zones and assign interfaces to each zone on primary logical system.

```
[edit security zones]
user@host#set security-zone root_trust interfaces reth1.0 host-inbound-traffic system-
services all
user@host#set security-zone root_trust interfaces reth1.0 host-inbound-traffic protocols all
user@host#set security-zone root_trust interfaces lt-0/0/0.1 host-inbound-traffic system-
services all
user@host#set security-zone root_trust interfaces lt-0/0/0.1 host-inbound-traffic protocols
all
user@host#set firewall family inet filter impair-ldap term allow_all then accept
```

Results

From configuration mode, confirm your configuration by entering the `show services user-identification identity-management` `show chassis cluster` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show services user-identification identity-management
connection {
    connect-method https;
    port 443;
    primary {
        address 192.0.2.5;
        client-id otest;
        client-secret "$ABC123"; ## SECRET-DATA
```

```
    }
}
```

```
user@host# show chassis cluster
reth-count 5;
    control-ports {
    fpc 3 port 0;
    fpc 9 port 0;
}
redundancy-group 0 {
    node 0 priority 200;
    node 1 priority 1;
}
redundancy-group 1 {
    node 0 priority 100;
    node 1 priority 1;
}
redundancy-group 2 {
    node 0 priority 100;
    node 1 priority 1;
}
redundancy-group 3 {
    node 0 priority 100;
    node 1 priority 1;
}
redundancy-group 4 {
    node 0 priority 100;
    node 1 priority 1;
}
```

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

Verification

IN THIS SECTION

- Verifying chassis cluster status and authentication entries | **203**
- Verifying chassis cluster status | **204**

To confirm that the configuration is working properly, perform the below tasks:

Verifying chassis cluster status and authentication entries

Purpose

To verify authentication entries in a logical system.

Action

To verify the configuration is working properly, enter the `show services user-identification authentication-table authentication-source identity-management logical-system all` command.

```
user@host> show services user-identification authentication-table authentication-source identity-
management logical-system all
node0:
--------------------------------------------------------------------------
Logical System: root-logical-system
Domain: ad2012.jims.com
Total entries: 3
Source IP       Username       groups(Ref by policy)          state
2001:db8:aaaa:  N/A                                           Valid
2001:db8:aaaa:  administrator                                 Valid
203.0.113.50    administrator                                 Valid
node1:
--------------------------------------------------------------------------
Logical System: root-logical-system
Domain: ad2012.jims.com
Total entries: 3
Source IP       Username       groups(Ref by policy)          state
2001:db8:aaaa:  N/A                                           Valid
2001:db8:aaaa:  administrator                                 Valid
203.0.113.50    administrator                                 Valid
```

Meaning

The output displays the authentication entries that are shared from user logical system to root logical system.

Verifying chassis cluster status

Purpose

Verify chassis cluster status after rebooting the primary node.

Action

To verify the configuration is working properly, enter the `show chassis cluster status` command.

```
user@host> show chassis cluster status
Monitor Failure codes:
CS  Cold Sync monitoring        FL  Fabric Connection monitoring
GR  GRES monitoring             HW  Hardware monitoring
IF  Interface monitoring        IP  IP monitoring
LB  Loopback monitoring         MB  Mbuf monitoring
NH  Nexthop monitoring          NP  NPC monitoring
SP  SPU monitoring              SM  Schedule monitoring
CF  Config Sync monitoring      RE  Relinquish monitoring
Cluster ID: 6
Node   Priority Status               Preempt Manual   Monitor-failures
Redundancy group: 0 , Failover count: 0
node0  200      hold                 no      no       None
node1  1        secondary            no      no       None
Redundancy group: 1 , Failover count: 0
node0  0        hold                 no      no       CS
node1  1        secondary            no      no       None
Redundancy group: 2 , Failover count: 0
node0  0        hold                 no      no       CS
node1  1        secondary            no      no       None
Redundancy group: 3 , Failover count: 0
node0  0        hold                 no      no       CS
node1  1        secondary            no      no       None
Redundancy group: 4 , Failover count: 0
node0  0        hold                 no      no       CS
node1  1        secondary            no      no       None
```

Meaning

The output displays user identification management session existing on lsys1 and lsys2 after rebooting the primary node.

SEE ALSO

show services user-identification authentication-table

Example: Configure Integrated User Firewall in Customized Model for Logical System

IN THIS SECTION

- Requirements | **205**
- Overview | **206**
- Configuration | **206**
- Verification | **210**

This example shows how to configure the integrated user firewall by using a customized model through the Juniper Identity Management Service (JIMS) server with active mode for a logical system. The primary logical systems does not share the authentication entries with the logical system. The SRX series device queries the authentication entries received from the JIMS server through HTTPs protocol in active mode.

In this example following configurations are performed:

- Active JIMS Server Configuration
- Logical System IP Query Configuration
- Logical System Authentication Entry Configuration
- Logical System Security Policy Configuration

Requirements

This example uses the following hardware and software components:

- JIMS server version 2.0
- Junos OS Release 19.3R1

Before you begin, be sure you have following information:

- The IP address of the JIMS server.
- The port number on the JIMS server for receiving HTTPs requests.
- The client ID from the JIMS server for active query server.
- The client secret from the JIMS server for active query server.

Overview

In this example, you can configure JIMS with HTTPs connection on port 443 and primary server with IPv4 address on the primary logical system, policy p2 with source-identity `group1` on logical system `LSYS1`.

Configuration

IN THIS SECTION

- CLI Quick Configuration | **206**
- Configuring Integrated User Firewall in Customized Model: | **207**
- Results | **208**

CLI Quick Configuration

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

```
set services user-identification logical-domain-identity-management active query-server jims1
connection connect-method https
set services user-identification logical-domain-identity-management active query-server jims1
connection port 443
set services user-identification logical-domain-identity-management active query-server jims1
connection primary address 192.0.2.5
set services user-identification logical-domain-identity-management active query-server jims1
connection primary client-id otest
set services user-identification logical-domain-identity-management active query-server jims1
connection primary client-secret "$ABC123"
set logical-systems LSYS1 services user-identification logical-domain-identity-management active
ip-query query-delay-time 30
set logical-systems LSYS1 services user-identification logical-domain-identity-management active
```

```
invalid-authentication-entry-timeout 1
set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy p2 match
source-address any
set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy p2 match
destination-address any
set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy p2 match
application any
set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy p2 match
source-identity "example.com\group1"
set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy p2 then permit
```

Configuring Integrated User Firewall in Customized Model:

Step-by-Step Procedure

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

To configuring Integrated User Firewall in Customized Model:

1. Configure JIMS as the authentication source for advanced query requests with the primary address. The SRX Series device requires this information to contact the server.

```
user@host# set services user-identification logical-domain-identity-management active query-
server jims1 connection connect-method https
user@host# set services user-identification logical-domain-identity-management active query-
server jims1 connection port 443
user@host# set services user-identification logical-domain-identity-management active query-
server jims1 connection primary address 192.0.2.5
user@host# set services user-identification logical-domain-identity-management active query-
server jims1 connection primary client-id otest
user@host# set services user-identification logical-domain-identity-management active query-
server jims1 connection primary client-secret "$ABC123"
```

2. Configure the IP query delay time for LSYS1.

```
user@host# set logical-systems LSYS1 services user-identification logical-domain-identity-
management active ip-query query-delay-time 30
```

3. Configure the authentication entry attributes for LSYS1.

```
user@host# set logical-systems LSYS1 services user-identification logical-domain-identity-
management active invalid-authentication-entry-timeout 1
```

4. Configure the security policy p2 that permits traffic from-zone untrust to-zone trust for LSYS1.

```
user@host# set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy
p2 match source-address any
user@host# set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy
p2 match destination-address any
user@host# set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy
p2 match application any
user@host# set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy
p2 match source-identity "example.com\group1"
user@host#set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy
p2 then permit
```

Results

From configuration mode, confirm your configuration by entering the `show services user-identification logical-domain-identity-management` and `show logical-systems LSYS1` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show services user-identification logical-domain-identity-management
active {
    query-server jims1 {
        connection {
            connect-method https;
            port 443;
            primary {
                address 192.0.2.5;
                client-id otest;
                client-secret "$ABC123"; ## SECRET-DATA
            }
        }
```

```
    }
}
```

```
user@host# show logical-systems LSYS1
security {
    policies {
        from-zone untrust to-zone trust {
            policy p2 {
                match {
                    source-address any;
                    destination-address any;
                    application any;
                    source-identity "example.com\group1";
                }
                then {
                    permit;
                }
            }
        }
    }
}
services {
    user-identification {
        logical-domain-identity-management {
            active {
                invalid-authentication-entry-timeout 1;
                ip-query {
                    query-delay-time 30;
                }
            }
        }
    }
}
```

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

Verification

IN THIS SECTION

- Verifying the User Identification Identity Management status | **210**
- Verifying the User Identification Identity Management status counters | **211**
- Verifying the User Identification Authentication Table | **212**

Verifying the User Identification Identity Management status

Purpose

Verify the user identification status for identity-management as the authentication source.

Action

To verify the configuration is working properly, enter the `show services user-identification logical-domain-identity-management status` command.

```
user@host> show services user-identification logical-domain-identity-management status
node0:
--------------------------------------------------------------------------
      Query server name                :jims1
      Primary server :
     Address                      : 192.0.2.5
     Port                         : 443
     Connection method            : HTTPS
     Connection status            : Online
     Last received status message : OK (200)
     Access token                 : isdHIbl8BXwxFftMRubGVsELRukYXtW3rtKmHiL
     Token expire time            : 2017-11-27 23:45:22
     Secondary server :
     Address                      : Not configured
```

Meaning

The output displays the statistical data about the advanced user query function batch queries and IP queries, or show status on the Juniper Identity Management Service servers.

Verifying the User Identification Identity Management status counters

Purpose

Verify the user identification counters for identity-management as the authentication source.

Action

To verify the configuration is working properly, enter the `show services user-identification logical-domain-identity-management counters` command.

```
user@host> show services user-identification logical-domain-identity-management counters
node0:
      --------------------------------------------------------------------------
   Query server name                      :jims1
    Primary server :
     Address                           : 192.0.2.5
     Batch query sent number           : 65381
     Batch query total response number : 64930
     Batch query error response number : 38
     Batch query last response time    : 2018-08-14 15:10:52
     IP query sent number              : 10
     IP query total response number    : 10
     IP query error response number    : 0
     IP query last response time       : 2018-08-13 12:41:56
  Secondary server :
 Address                           : Not configured
```

Meaning

The output displays the statistical data about the advanced user query function batch queries and IP queries, or show counters on the Juniper Identity Management Service servers.

Verifying the User Identification Authentication Table

Purpose

Verify the user identity information authentication table entries for the specified authentication source.

Action

To verify the configuration is working properly, enter the `show services user-identification authentication-table authentication-source all logical-system LSYS1` command.

```
user@host> show services user-identification authentication-table authentication-source all
logical-system LSYS1
node0:
      --------------------------------------------------------------------------
   Logical System: LSYS1
Domain: example.com
Total entries: 4
Source IP           Username       groups(Ref by policy)          state
10.12.0.2         administrator  posture-healthy               Valid
10.12.0.15        administrator  posture-healthy               Valid
2001:db8::5         N/A            posture-healthy                Valid
2001:db8::342c:302b N/A            posture-healthy                Valid
```

Meaning

The output displays the entire content of the specified authentication source's authentication table, or a specific domain, group, or user based on the user name. Display the identity information for a user based on the IP address of the user's device.

Release History Table

Release	Description
19.3R1	Starting in Junos OS Release 19.3R1, support for user firewall authentication is enhanced by using a customized model through integrated JIMS with active mode.
18.3R1	Starting in Junos OS Release 18.3R1, the support for authentication sources is extended to include Local authentication, Active Directory (AD) authentication, and firewall authentication in addition to the existing support for authentication sources Juniper Identity Management Service (JIMS) and ClearPass authentication.

18.2R1	Starting in Junos OS Release 18.2R1, the support for user firewall authentication is enhanced using a shared model. In this model, user logical systems share user firewall configuration and authentication entries with the primary logical system and the integrated user firewall authentication is supported in a user logical system.

RELATED DOCUMENTATION

Example: Configuring Security log stream for Logical Systems | 96

Security Policies for Logical Systems

IN THIS SECTION

- Understanding Logical Systems Security Policies | 213
- Example: Configuring Security Policies in a User Logical Systems | 216
- Configuring Dynamic Address for Logical Systems | 221

Security policies are used to secure business and control access to LAN resources. Secure access is required both within the company across the LAN and in its interactions with external networks such as the Internet. Junos OS provides powerful network security features through its stateful firewall, application firewall, and user identity firewall. All three types of firewall enforcement are implemented through security policies. For more information, see the following topics:

Understanding Logical Systems Security Policies

IN THIS SECTION

- Security Policies in Logical Systems | 214
- Application Timeouts | 215

Security Policy Allocation | 215

Security Policies in Logical Systems

Security policies enforce rules for what traffic can pass through the firewall and actions that need to take place on the traffic as it passes through the firewall. From the perspective of security policies, traffic enters one security zone and exits another security zone.

By default, a logical system denies all traffic in all directions, including intra-zone and inter-zone directions. Through the creation of security policies, the logical system administrator can control the traffic flow from zone to zone by defining the kinds of traffic permitted to pass from specified sources to specified destinations.

Security policies can be configured in the primary logical system and in user logical systems. Configuring a security policy in a logical system is the same as configuring a security policy on a device that is not configured for logical systems. Any security policies, policy rules, address books, applications and application sets, and schedulers created within a logical system are only applicable to that logical system. Only predefined applications and application sets, such as `junos-ftp`, can be shared between logical systems.

NOTE: In a logical system, you cannot specify `global` as either the from-zone or the to-zone in a security policy.

The user logical system administrator can configure and view all attributes for security policies in a user logical system. All attributes of a security policy in a user logical system are also visible to the primary administrator.

Starting in Junos OS Release 18.4R1, the user can create dynamic address within a logical system. A dynamic address entry contains IP addresses and prefixes extracted from external sources. The security policies use the dynamic address in the source-address field or destination-address field.

A dynamic address entry (DAE) is a group of IP addresses that can be entered manually or imported from external sources within logical systems. The DAE feature allows feed-based IP objects to be used in security policies to either deny or allow traffic based on either source or destination IP criteria.

NOTE: The maximum number of DAE depends on the dynamic-addresses assigned to the logical systems. Starting in Junos 18.4R1, the `set security dynamic-address feed-server` command can be configured under the logical systems.

Application Timeouts

The application timeout value set for an application determines the session timeout. Application timeout behavior is the same in a logical system as at the root level. However, user logical system administrators can use predefined applications in security policies but cannot modify the timeout value of predefined applications. This is because the predefined applications are shared by the primary logical system and all user logical systems, so the user logical system administrator is not allowed to change its behavior. Application timeout values are stored in the application entry database and in the corresponding logical system TCP and UDP port-based timeout tables.

If the application that is matched for the traffic has a timeout value, that timeout value is used. Otherwise, the lookup proceeds in the following order until an application timeout value is found:

1. The logical system TCP and UDP port-based timeout table is searched for a timeout value.
2. The root TCP and UDP port-based timeout table is searched for a timeout value.
3. The protocol-based default timeout table is searched for a timeout value.

Security Policy Allocation

The primary administrator configures the maximum and reserved numbers of security policies for each user logical system. The user logical system administrator can then create security policies in the user logical system. From a user logical system, the user logical system administrator can use the `show system security-profile policy` command to view the number of security policies allocated to the user logical system.

NOTE: The primary administrator can configure a security profile for the primary logical system that specifies the maximum and reserved numbers of security policies applied to the primary logical system. The number of policies configured in the primary logical system count toward the maximum number of policies available on the device.

SEE ALSO

Example: Configuring Security Policies in a User Logical Systems | 216
Understanding Logical Systems Security Profiles (Primary Administrators Only) | 71
User Logical Systems Configuration Overview | 51
Security Policies Overview
Understanding Policy Application Timeout Configuration and Lookup

Example: Configuring Security Policies in a User Logical Systems

IN THIS SECTION

- Requirements | 216
- Overview | 217
- Configuration | 218
- Verification | 220

This example shows how to configure security policies for a user logical system.

Requirements

Before you begin:

- Log in to the user logical system as the logical system administrator. See "User Logical Systems Configuration Overview" on page 51.
- Use the `show system security-profiles policy` command to see the security policy resources allocated to the logical system.
- Configure zones and address books. See "Example: Configuring Security Zones for a User Logical Systems" on page 167.

Overview

IN THIS SECTION

- Topology | 217

This example configures the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System" on page 57.

This example configures the security policies described in Table 17 on page 217.

Table 17: User Logical System Security Policies Configuration

Name	Configuration Parameters
permit-all-to-otherlsys	Permit the following traffic: • From zone: ls-product-design-trust • To zone: ls-product-design-untrust • Source address: product-designers • Destination address: otherlsys • Application: any
permit-all-from-otherlsys	Permit the following traffic: • From zone: ls-product-design-untrust • To zone: ls-product-design-trust • Source address: otherlsys • Destination address: product-designers • Application: any

Topology

Configuration

IN THIS SECTION

- Procedure | 218

Procedure

CLI Quick Configuration

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

```
set security policies from-zone ls-product-design-trust to-zone  ls-product-design-untrust
policy permit-all-to-otherlsys match source-address product-designers
set security policies from-zone ls-product-design-trust to-zone  ls-product-design-untrust
policy permit-all-to-otherlsys match destination-address otherlsys
set security policies from-zone ls-product-design-trust to-zone  ls-product-design-untrust
policy permit-all-to-otherlsys match application any
set security policies from-zone ls-product-design-trust to-zone  ls-product-design-untrust
policy permit-all-to-otherlsys then permit
set security policies from-zone ls-product-design-untrust to-zone  ls-product-design-trust
policy permit-all-from-otherlsys match source-address otherlsys
set security policies from-zone ls-product-design-untrust to-zone  ls-product-design-trust
policy permit-all-from-otherlsys match destination-address product-designers
set security policies from-zone ls-product-design-untrust to-zone  ls-product-design-trust
policy permit-all-from-otherlsys match application any
set security policies from-zone ls-product-design-untrust to-zone  ls-product-design-trust
policy permit-all-from-otherlsys then permit
```

Step-by-Step Procedure

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

To configure security policies in a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

```
lsdesignadmin1@host:ls-product-design> configure
lsdesignadmin1@host:ls-product-design#
```

2. Configure a security policy that permits traffic from the ls-product-design-trust zone to the ls-product-design-untrust zone.

```
[edit security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust]
lsdesignadmin1@host:ls-product-design# set policy permit-all-to-otherlsys match source-address product-designers
lsdesignadmin1@host:ls-product-design# set policy permit-all-to-otherlsys match destination-address otherlsys
lsdesignadmin1@host:ls-product-design# set policy permit-all-to-otherlsys match application any
lsdesignadmin1@host:ls-product-design# set policy permit-all-to-otherlsys then permit
```

3. Configure a security policy that permits traffic from the ls-product-design-untrust zone to the ls-product-design-trust zone.

```
[edit security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust]
lsdesignadmin1@host:ls-product-design# set policy permit-all-from-otherlsys match source-address otherlsys
lsdesignadmin1@host:ls-product-design# set policy permit-all-from-otherlsys match destination-address product-designers
lsdesignadmin1@host:ls-product-design# set policy permit-all-from-otherlsys match application any
lsdesignadmin1@host:ls-product-design# set policy permit-all-from-otherlsys then permit
```

Results

From configuration mode, confirm your configuration by entering the `show security policies` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
lsdesignadmin1@host:ls-product-design# show security policies
from-zone ls-product-design-trust to-zone ls-product-design-untrust {
    policy permit-all-to-otherlsys {
        match {
```

```
            source-address product-designers;
            destination-address otherlsys;
            application any;
        }
        then {
            permit;
        }
    }
}
    from-zone ls-product-design-untrust to-zone ls-product-design-trust {
        policy permit-all-from-otherlsys {
            match {
                source-address otherlsys;
                destination-address product-designers;
                application any;
            }
            then {
                permit;
            }
        }
    }
```

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

Verification

IN THIS SECTION

- Verifying Policy Configuration | 220

To confirm that the configuration is working properly, perform these tasks:

Verifying Policy Configuration

Purpose

Verify information about policies and rules.

Action

From operational mode, enter the `show security policies detail` command to display a summary of all policies configured on the logical system.

SEE ALSO

Understanding Logical Systems Security Policies | 213

User Logical Systems Configuration Overview | 51

Troubleshooting Security Policies

Configuring Dynamic Address for Logical Systems

A dynamic address entry in logical systems provides dynamic IP address information to security policies. To use dynamic address, you must specify basic information of dynamic address including their names, feeds and properties for a logical system.

- Read the Example: Configuring Security Policies in a User Logical Systems to understand how and where this procedure fits to configure the security policy.

To configure dynamic address in IPv4 networks within a logical system:

1. Define the logical system name as LSYS1.

```
[edit]
user@host# set logical-systems LSYS1
```

2. Create dynamic address within a logical system.

```
[edit logical-systems LSYS1]
user@host# set security dynamic-address address-name Ipv4 profile category IPFilter feed fd1
```

3. Confirm your configuration by entering the `show logical-systems LSYS1 security dynamic-address` command.

```
[edit]
user@host# show logical-systems LSYS1 security dynamic-address
address-name Ipv4 {
    profile {
```

```
        category GeoIP;
        category IPFilter {
            feed fd1;
            }
        }
    }
}
```

- To configure the security policies in the logical system:

  1. Define the logical system name as LSYS1.

     ```
     [edit]
     user@host# set logical-systems LSYS1
     ```

  2. Create a security policy as p1 that permits traffic from zone trust to zone untrust and configure the match condition.

     ```
     [edit logical-systems LSYS1 security policies from-zone trust to-zone untrust]
     user@host# set policy p1 match source-address any
     user@host# set policy p1 match destination-address any
     user@host# set policy p1 match application any
     user@host# set policy p1 then permit
     ```

  3. Confirm your configuration by entering the `show logical-systems LSYS1 security policies` command.

     ```
     [edit]
     user@host# show logical-systems LSYS1 security policies
     from-zone trust to-zone untrust {
         policy p1 {
             match {
                 source-address any;
                 destination-address any;
                 application any;
             }
             then {
                 permit;
             }
     ```

```
    }
}
```

Screen Options for User Logical Systems

IN THIS SECTION

- Understanding Logical Systems Screen Options | 223
- Example: Configuring Screen Options for a User Logical Systems | 224

Screen options on SRX Series devices prevent attacks, such as IP address sweeps, port scans, denial of service (DOS) attacks, ICMP, UDP, and SYN floods. For more information, see the following topics:

Understanding Logical Systems Screen Options

Junos OS screen options secure a zone by inspecting, then allowing or denying, all connection attempts that require crossing an interface bound to that zone. Junos OS then applies firewall policies, which can contain content filtering and IDP components, to the traffic that passes the screen filters.

All screen options available on the device are available in each logical system. Each user logical system administrator can configure screen options for their user logical system. The primary administrator can configure screen options for the primary logical system as well as all user logical systems.

The user logical system administrator can configure and view all screen options in a user logical system. All screen options in a user logical system are visible to the primary administrator.

SEE ALSO

Example: Configuring Screen Options for a User Logical Systems | 224

User Logical Systems Configuration Overview | 51

Attack Detection and Prevention Overview

Example: Configuring Screen Options for a User Logical Systems

IN THIS SECTION

- Requirements | 224
- Overview | 224
- Configuration | 225

This example shows how to configure screen options for a user logical system.

Requirements

Before you begin:

- Log in to the user logical system as the user logical system administrator. See "User Logical Systems Configuration Overview" on page 51.
- Configure zones for the user logical system. See "Example: Configuring Security Zones for a User Logical Systems" on page 167.

Overview

This example configures the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System" on page 57.

You can limit the number of concurrent sessions to the same destination IP address in a user logical system. Setting a destination-based session limit can ensure that Junos OS allows only an acceptable number of concurrent connection requests—no matter what the source—to reach any one host. When the number of concurrent connection requests to an IP address surpasses the limit, Junos OS blocks further connection attempts to that IP address. This example creates the screen options described in Table 18 on page 224.

Table 18: User Logical System Screen Options Configuration

Name	Configuration Parameters
limit-destination-sessions	• Limits concurrent connection requests to destination IPs to 80. • Applied to ls-product-design-untrust zone.

Configuration

IN THIS SECTION

- Procedure | 225

Procedure

CLI Quick Configuration

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

```
set security screen ids-option limit-destination-sessions limit-session destination-ip-based 80
set security zones security-zone ls-product-design-untrust screen limit-destination-sessions
```

Step-by-Step Procedure

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

To configure destination-based session limits in a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

```
lsdesignadmin1@host:ls-product-design> configure
lsdesignadmin1@host:ls-product-design#
```

2. Configure a screen option for a destination-based session limit.

```
[edit security]
lsdesignadmin1@host:ls-product-design# set screen ids-option limit-destination-sessions limit-
session destination-ip-based 80
```

3. Set the security zone for the screen option.

```
[edit security]
lsdesignadmin1@host:ls-product-design#  set zones security-zone ls-product-design-untrust
screen limit-destination-sessions
```

Results

From configuration mode, confirm your configuration by entering the `show security screen` and `show security zone` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this `show` command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
lsdesignadmin1@host:ls-product-design# show security screen
ids-option limit-destination-sessions {
    limit-session {
        destination-ip-based 80;
    }
}
lsdesignadmin1@host:ls-product-design# show security zones
security-zone ls-product-design-trust {
    ...
}
    security-zone ls-product-design-untrust {
        screen limit-destination-sessions;
        ...
    }
```

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

SEE ALSO

User Logical Systems Configuration Overview | 51

Understanding Logical Systems Screen Options | 223

RELATED DOCUMENTATION

IDP for Logical Systems | 273

Secure Wire for Logical Systems

IN THIS SECTION

- Secure Wire for Logical Systems Overview | 227
- Example: Configure Secure Wire for User Logical Systems | 229

Secure Wire for Logical Systems Overview

IN THIS SECTION

- Limitations | 229

You can forward the traffic that arrives on a specific interface without any change through another interface on logical systems. This mapping of interfaces on logical systems is called secure wire. Secure wire allows an SRX Series device to deploy in the path of network traffic without changing the routing tables or a reconfiguration of neighboring devices. Figure 7 on page 228 shows a typical in-path deployment of an SRX Series device with secure wire.

Figure 7: SRX Series Device In-Path Deployment with Secure Wire

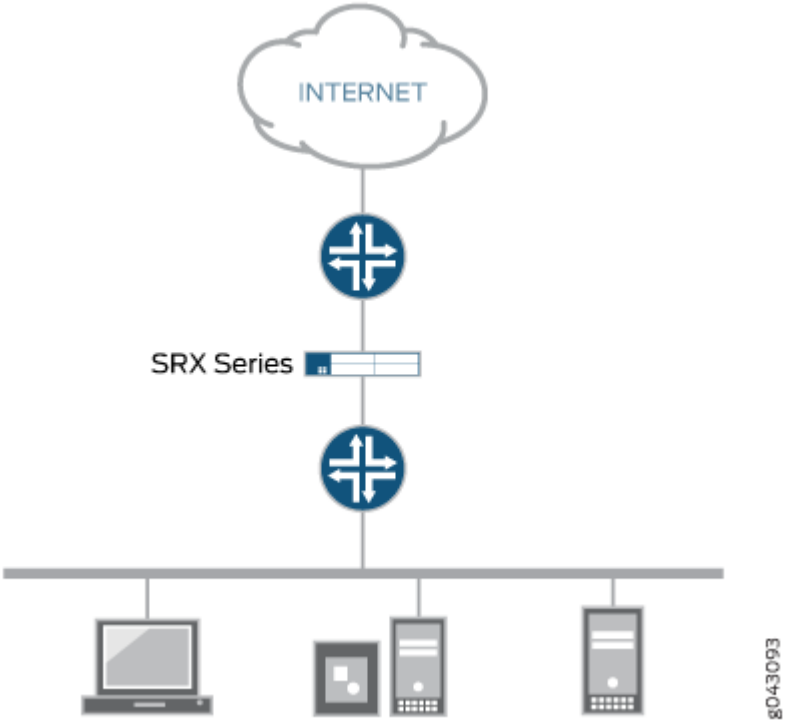

Secure wire maps two peer interfaces. It differs from transparent and route modes, and there is no switching or routing lookup to forward traffic. When security policy permits the traffic, secure wire forwards a packet arriving on one peer interface immediately to the other peer interface without change. There is no routing or switching decision made on the packet. Secure wire also forwards the return traffic unchanged. The secure wire feature is supported for both IPv4 and IPv6 traffic on Ethernet logical interfaces only.

Secure wire is a special case of Layer 2 transparent mode on SRX Series devices that provide point-to-point connections. This means that the two interfaces of a secure wire must directly connect to Layer 3 entities, such as routers or hosts. You can connect secure wire interfaces to switches. However, note that when security policy permits traffic, a secure wire interface forwards all arriving traffic to the peer interface.

Secure wire can coexist with Layer 3 mode. While you configure Layer 2 and Layer 3 interfaces at the same time, traffic forwarding occurs independently on Layer 2 and Layer 3 interfaces.

Secure wire can coexist with Layer 2 transparent mode. If both features exist on the same SRX Series device, you need to configure them in different VLANs.

Secure wire support for root logical system extends to user logical systems. You can forward traffic immediately that arrives on a specific interface to another interface without modifying any received frames on the user logical systems.

Limitations

Secure wire doesn't support:

- IRB interface
- Z-mode
- MPLS label encapsulation
- Tenant system
- Interconnect logical system

Example: Configure Secure Wire for User Logical Systems

IN THIS SECTION

- Requirements | 229
- Overview | 230
- Configuration | 230
- Verification | 231

In this example, you can configure secure wire for a user logical system and forward traffic from one interface to another interface without changing any frame.

Requirements

Before you begin:

- Configure security profile for a user logical system, see "Example: Configuring User Logical Systems Security Profiles" on page 89.

Overview

In this example, you can configure 10-Gigabit Ethernet interfaces xe-1/0/1 and xe-1/0/2 under a user logical system, called LSYS1. You can configure secure wire resource allocation per logical system. When traffic passes to xe-1/0/1 interface, without changing any frame, secure wire forwards the traffic to xe-1/0/2 interface based on the defined security policy.

Configuration

IN THIS SECTION

- Procedure | 230
- Results | 231

Procedure

CLI Quick Configuration

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

```
user@host#set logical-systems LSYS1 security forwarding-options secure-wire myLSYS1sw01
interface xe-1/0/1.0
user@host#set logical-systems LSYS1 security forwarding-options secure-wire myLSYS1sw01
interface xe-1/0/2.0
user@host#set system security-profile prof1 secure-wire maximum 100
user@host#set system security-profile prof1 secure-wire reserved 1
```

Step-by-Step Procedure

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

1. Configure secure wire under a user logical system.

```
[edit]
user@host#set logical-systems LSYS1 security forwarding-options secure-wire myLSYS1sw01
```

```
interface xe-1/0/1.0
user@host#set logical-systems LSYS1 security forwarding-options secure-wire myLSYS1sw01
interface xe-1/0/2.0
```

2. Create the security profile, and specify the number of maximum and reserved quota.

```
[edit]
user@host#set system security-profile prof1 secure-wire maximum 100
user@host#set system security-profile prof1 secure-wire reserved 1
```

Results

From configuration mode, confirm your configuration by entering the `show logical-systems LSYS1 security forwarding-options secure-wire myLSYS1sw01`, and `show system security-profile prof1` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host#show logical-systems LSYS1 security forwarding-options secure-wire myLSYS1sw01
interface [ xe-1/0/1.0 xe-1/0/2.0 ];
```

```
user@host#show system security-profile prof1
    secure-wire {
        maximum 100;
        reserved 1;
    }
    logical-system LSYS1;
```

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

Verification

IN THIS SECTION

- Verify Secure Wire Mapping | 232
- Verify Resource Allocation | 232

Confirm that the configuration is working properly.

Verify Secure Wire Mapping

Purpose

Verify the secure wire mapping.

Action

From operational mode, enter the `show security forward-options secure-wire logical-system LSYS1` command.

```
Logical System       Secure wire                    Interface       Link  Interface       Link

LSYS1                myLSYS1sw01                    xe-1/0/1.0      up    xe-1/0/2.0      up
Total secure wires: 1
```

Verify Resource Allocation

Purpose

Verify the resource allocation for a user logical system.

Action

From operational mode, enter the `show system security-profile secure-wire logical-system LSYS1` command.

```
logical-system tenant name   security profile name       usage    reserved     maximum

LSYS1                        prof1                           1           1         100
```

RELATED DOCUMENTATION

security-profile | 871

secure-wire

show security forward-options secure-wire
show system security-profile secure-wire | 1305

VPNs in Logical Systems

IN THIS SECTION

- Understanding Route-Based VPN Tunnels in Logical Systems | 233
- Example: Configuring IKE and IPsec SAs for a VPN Tunnel (Primary Administrators Only) | 235
- Example: Configuring a Route-Based VPN Tunnel in a User Logical Systems | 245

A VPN is an encrypted connection over the Internet from a device to a network. The encrypted connection helps ensure that sensitive data is safely transmitted. VPN prevents unauthorized access eavesdropping on the traffic, and allows the user to conduct work remotely. For more information, see the following topics:

Understanding Route-Based VPN Tunnels in Logical Systems

A VPN connection can secure traffic that passes between a logical system and a remote site across a WAN. With route-based VPNs, you configure one or more security policies in a logical system to regulate the traffic flowing through a single IP Security (IPsec) tunnel. For each IPsec tunnel, there is one set of IKE and IPsec security associations (SAs) that must be configured at the root level by the primary administrator.

NOTE: The external interface configured under the gateway configuration can only be a part of the root logical system.

NOTE: Only route-based VPNs are supported in logical systems. Policy-based VPNs are not supported.

In addition to configuring IKE and IPsec SAs for each VPN, the primary administrator must also assign a secure tunnel (st0) interface to a user logical system. An st0 interface can only be assigned to a single user logical system. However, multiple user logical systems can each be assigned their own st0 interface.

NOTE: The st0 unit 0 interface should not be assigned to a logical system, as an SA cannot be set up for this interface.

The user logical system administrator can configure the IP address and other attributes of the st0 interface assigned to the user logical system. The user logical system administrator cannot delete an st0 interface assigned to their user logical system.

For route-based VPNs, a security policy refers to a destination address and not a specific VPN tunnel. For cleartext traffic in a user logical system to be sent to the VPN tunnel for encapsulation, the user logical system administrator must make the following configurations:

- Security policy that permits traffic to a specified destination.
- Static route to the destination with the st0 interface as the next hop.

When Junos OS looks up routes in the user logical system to find the interface to use to send traffic to the destination address, it finds a static route through the st0 interface. Traffic is routed to the VPN tunnel as long as the security policy action is permit.

NOTE: Traffic selectors are not supported in logical systems.

The primary logical system and a user logical system can share a route-based VPN tunnel. An st0 interface assigned to a user logical system can also be used by the primary logical system. For the primary logical system, the primary administrator configures a security policy that permits traffic to the remote destination and a static route to the remote destination with the st0 interface as the next hop.

VPN monitoring is configured by the primary administrator in the primary logical system. For the VPN monitor source interface, the primary administrator must specify the st0 interface; a physical interface for a user logical system cannot be specified.

SEE ALSO

Understanding Route-Based IPsec VPNs

User Logical Systems Configuration Overview | **51**

Example: Configuring IKE and IPsec SAs for a VPN Tunnel (Primary Administrators Only) | **235**

Example: Configuring a Route-Based VPN Tunnel in a User Logical Systems | **245**

Example: Configuring IKE and IPsec SAs for a VPN Tunnel (Primary Administrators Only)

IN THIS SECTION

- Requirements | 235
- Overview | 235
- Configuration | 238
- Verification | 242

The primary administrator is responsible for assigning an st0 interface to a user logical system and configuring IKE and IPsec SAs at the root level for each VPN tunnel. This example shows how to assign an st0 interface to a user logical system and configure IKE and IPsec SA parameters.

Requirements

Before you begin:

- Log in to the primary logical system as the primary administrator. See "'Understanding the Primary Logical Systems and the Primary Administrator Role" on page 23.
- Read *Understanding Route-Based IPsec VPNs*.

Overview

IN THIS SECTION

- Topology | 237

In this example you configure a VPN tunnel for the ls-product-design user logical system. This example configures the VPN tunnel parameters described in Table 19 on page 236.

Table 19: Logical System VPN Tunnel Configuration

Feature	Name	Configuration Parameters
Tunnel interface	st0 unit 1	Assigned to ls-product-design logical system
IKE proposal	ike-phase1-proposal	• Preshared keys authentication • Diffie-Hellman group 2 • sha1 authentication algorithm • aes-128-cbc encryption algorithm
IKE policy		• Main mode • References IKE proposal ike-phase1-proposal • ASCII preshared key 395psksecr3t
IKE gateway	ike-gw	• External interface ge-0/0/3.0 • References IKE policy ike-phase1-policy • Address 2.2.2.2
IPsec proposal	ipsec-phase2-proposal	• ESP protocol • hmac-sha1-96 authentication algorithm • aes-128-cbc encryption algorithm
IPsec policy	vpn-policy1	• References ipsec-phase2-proposal • perfect-forward-secrecy keys group2

Table 19: Logical System VPN Tunnel Configuration *(Continued)*

Feature	Name	Configuration Parameters
VPN	ike-vpn	• bind-interface st0.1 • References ike-gw gateway • References vpn-policy1 policy
VPN monitoring		For ike-vpn VPN: • source-interface st0.1 • destination-ip 4.0.0.1

Topology

Figure 8 on page 237 shows the topology for logical systems VPN tunnel.

Figure 8: Logical systems VPN tunnel

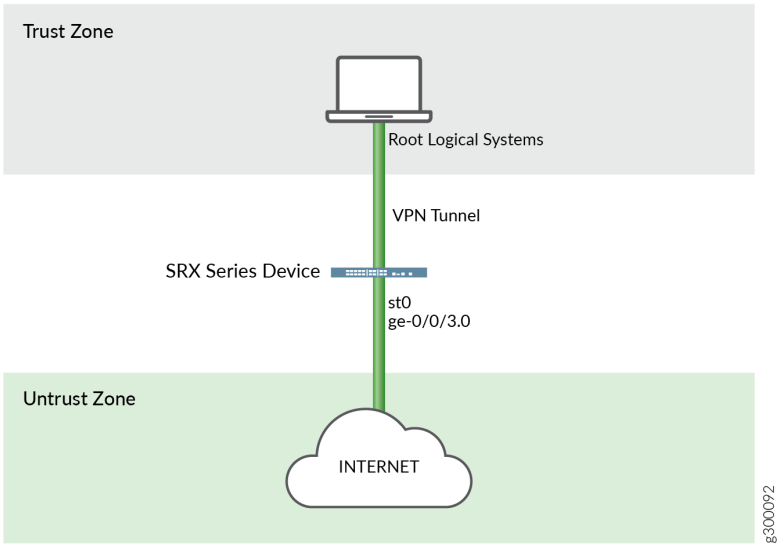

Configuration

IN THIS SECTION

- Procedure | 238

Procedure

CLI Quick Configuration

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

```
set logical-systems ls-product-design interfaces st0 unit 1
set security ike proposal ike-phase1-proposal authentication-method pre-shared-keys
set security ike proposal ike-phase1-proposal dh-group group2
set security ike proposal ike-phase1-proposal authentication-algorithm sha1
set security ike proposal ike-phase1-proposal encryption-algorithm aes-128-cbc
set security ike policy ike-phase1-policy mode main
set security ike policy ike-phase1-policy proposals ike-phase1-proposal
set security ike policy ike-phase1-policy pre-shared-key ascii-text "$ABC123"
set security ike gateway ike-gw ike-policy ike-phase1-policy
set security ike gateway ike-gw address 2.2.2.2
set security ike gateway ike-gw external-interface ge-0/0/3.0
set security ipsec proposal ipsec-phase2-proposal protocol esp
set security ipsec proposal ipsec-phase2-proposal authentication-algorithm hmac-sha1-96
set security ipsec proposal ipsec-phase2-proposal encryption-algorithm aes-128-cbc
set security ipsec policy vpn-policy1 perfect-forward-secrecy keys group2
set security ipsec policy vpn-policy1 proposals ipsec-phase2-proposal
set security ipsec vpn ike-vpn bind-interface st0.1
set security ipsec vpn ike-vpn vpn-monitor source-interface st0.1
set security ipsec vpn ike-vpn vpn-monitor destination-ip 4.0.0.1
set security ipsec vpn ike-vpn ike gateway ike-gw
set security ipsec vpn ike-vpn ike ipsec-policy vpn-policy1
```

Step-by-Step Procedure

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

To assign a VPN tunnel interface to a user logical system and configure IKE and IPsec SAs:

1. Log in to the primary logical system as the primary administrator and enter configuration mode.

```
[edit]
admin@host> configure
admin@host#
```

2. Assign a VPN tunnel interface.

```
[edit logical-systems ls-product-design]
admin@host# set interfaces st0 unit 1
```

3. Configure an IKE proposal.

```
[edit security ike]
admin@host# set proposal ike-phase1-proposal authentication-method pre-shared-keys
admin@host# set proposal ike-phase1-proposal dh-group group2
admin@host# set proposal ike-phase1-proposal authentication-algorithm sha1
admin@host# set proposal ike-phase1-proposal encryption-algorithm aes-128-cbc
```

4. Configure an IKE policy.

```
[edit security ike]
admin@host# set policy ike-phase1-policy mode main
admin@host# set policy ike-phase1-policy proposals ike-phase1-proposal
admin@host# set policy ike-phase1-policy pre-shared-key ascii-text 395psksecr3t
```

5. Configure an IKE gateway.

```
[edit security ike]
admin@host# set gateway ike-gw external-interface ge-0/0/3.0
```

```
admin@host# set gateway ike-gw ike-policy ike-phase1-policy
admin@host# set gateway ike-gw address 2.2.2.2
```

6. Configure an IPsec proposal.

```
[edit security ipsec]
admin@host# set proposal ipsec-phase2-proposal protocol esp
admin@host# set proposal ipsec-phase2-proposal authentication-algorithm hmac-sha1-96
admin@host# set proposal ipsec-phase2-proposal encryption-algorithm aes-128-cbc
```

7. Configure an IPsec policy.

```
[edit security ipsec]
admin@host# set policy vpn-policy1 proposals ipsec-phase2-proposal
admin@host# set policy vpn-policy1 perfect-forward-secrecy keys group2
```

8. Configure the VPN.

```
[edit security ipsec]
admin@host# set vpn ike-vpn bind-interface st0.1
admin@host# set vpn ike-vpn ike gateway ike-gw
admin@host# set vpn ike-vpn ike ipsec-policy vpn-policy1
```

9. Configure VPN monitoring.

```
[edit security ipsec]
admin@host# set vpn ike-vpn vpn-monitor source-interface st0.1
admin@host# set vpn ike-vpn vpn-monitor destination-ip 4.0.0.1
```

Results

From configuration mode, confirm your configuration by entering the `show interfaces`, `show security ike`, and `show security ipsec` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
admin@host# show interfaces
```

```
    st0 {
        unit 1;
    }
[edit]
admin@host# show security ike
    proposal ike-phase1-proposal {
        authentication-method pre-shared-keys;
        dh-group group2;
        authentication-algorithm sha1;
        encryption-algorithm aes-128-cbc;
    }
        policy ike-phase1-policy {
            mode main;
            proposals ike-phase1-proposal;
            pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
        }
        gateway ike-gw {
            ike-policy ike-phase1-policy;
            address 2.2.2.2;
            external-interface ge-0/0/3.0;
        }
[edit]
admin@host# show security ipsec
    proposal ipsec-phase2-proposal {
        protocol esp;
        authentication-algorithm hmac-sha1-96;
        encryption-algorithm aes-128-cbc;
    }
        policy vpn-policy1 {
            perfect-forward-secrecy {
                keys group2;
            }
            proposals ipsec-phase2-proposal;
        }
        vpn ike-vpn {
            bind-interface st0.1;
            vpn-monitor {
                source-interface st0.1;
                destination-ip 4.0.0.1;
            }
            ike {
                gateway ike-gw;
                ipsec-policy vpn-policy1;
```

```
        }
    }
```

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

Verification

IN THIS SECTION

- Verifying the IKE on Logical System | **242**
- Verifying the IPsec on Logical System | **243**

To confirm that the configuration is working properly, perform these tasks:

Verifying the IKE on Logical System

Purpose

Verify that the IKE support on Logical Systems.

Action

From operational mode, enter the `show security ike sa detail` command.

```
user@host> show security ike sa detail
IKE peer 2.2.2.2, Index 7796166, Gateway Name: GW1
  Role: Initiator, State: UP
  Initiator cookie: a1a6b1516bc43d54, Responder cookie: f0846e4239c817f8
  Exchange type: Aggressive, Authentication method: Pre-shared-keys
  Local: 3.3.3.2:500, Remote: 2.2.2.2:500
  Lifetime: Expires in 3585 seconds
  Reauth Lifetime: Disabled
  IKE Fragmentation: Disabled, Size: 0
  Remote Access Client Info: Unknown Client
  Peer ike-id: 2.2.2.2
  AAA assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha256-128
```

```
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha256
   Diffie-Hellman group  : DH-group-14
  Traffic statistics:
   Input  bytes  :                1056
   Output bytes  :                1311
   Input  packets:                   2
   Output packets:                   4
   Input  fragmentated packets:      0
   Output fragmentated packets:      0
  IPSec security associations: 1 created, 0 deleted
  Phase 2 negotiations in progress: 1

    Negotiation type: Quick mode, Role: Initiator, Message ID: 0
    Local: 3.3.3.2:500, Remote: 2.2.2.2:500
    Local identity: r0r2_store1@juniper.net
    Remote identity: 2.2.2.2
    Flags: IKE SA is created
```

Meaning

The output displays summary information about ike details.

Verifying the IPsec on Logical System

Purpose

Verify that the IPsec SA support on Logical Systems.

Action

From operational mode, enter the `show security ipsec sa detail` command.

```
user@host> show security ipsec sa detail
ID: 67109793 Virtual-system: root, VPN Name: VPN1
  Local Gateway: 3.3.3.2, Remote Gateway: 2.2.2.2
  Traffic Selector Name: VPN1_TS1
  Local Identity: ipv4(51.0.1.0-51.0.1.255)
  Remote Identity: ipv4(41.0.1.0-41.0.1.255)
  Version: IKEv1
  DF-bit: clear, Copy-Outer-DSCP Disabled, Bind-interface: st0.1
```

```
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0x2c608b29
  Tunnel events:
    Wed Aug 16 2017 23:50:07 -0700: IPSec SA negotiation successfully completed (1 times)
    Wed Aug 16 2017 23:50:07 -0700: IKE SA negotiation successfully completed (1 times)
    Wed Aug 16 2017 23:49:46 -0700: Negotiation failed with error code AUTHENTICATION_FAILED
received from peer (2 times)
    Wed Aug 16 2017 23:49:30 -0700: Tunnel is ready. Waiting for trigger event or peer to
trigger negotiation (1 times)
  Direction: inbound, SPI: e651d79e, AUX-SPI: 0, VPN Monitoring: -
    Hard lifetime: Expires in 2552 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1988 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha256-128, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
  Direction: outbound, SPI: 8ac9ce8, AUX-SPI: 0, VPN Monitoring: -
    Hard lifetime: Expires in 2552 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1988 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha256-128, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
```

Meaning

The output displays summary information about ipsec details.

SEE ALSO

Example: Configuring a Route-Based VPN Tunnel in a User Logical Systems | 245

Understanding Route-Based VPN Tunnels in Logical Systems | 233

User Logical Systems Configuration Overview | 51

Example: Configuring a Route-Based VPN Tunnel in a User Logical Systems

IN THIS SECTION

- Requirements | 245
- Overview | 245
- Configuration | 246
- Verification | 249

This example shows how to configure a route-based VPN tunnel in a user logical system.

Requirements

Before you begin:

- Log in to the user logical system as the logical system administrator. See "User Logical Systems Configuration Overview" on page 51.
- Ensure that an st0 interface is assigned to the user logical system and IKE and IPsec SAs are configured at the root level by the primary administrator. See "Example: Configuring IKE and IPsec SAs for a VPN Tunnel (Primary Administrators Only)" on page 235.

Overview

In this example, you configure the ls-product-design user logical system as shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System" on page 57.

You configure the route-based VPN parameters described in Table 20 on page 245.

Table 20: User Logical System Route-Based VPN Configuration

Feature	Name	Configuration Parameters
Tunnel interface	st0 unit 1	• IPv4 protocol family (inet) • IP address 10.11.11.150/24

Table 20: User Logical System Route-Based VPN Configuration *(Continued)*

Feature	Name	Configuration Parameters
Static route		• Destination 192.168.168.0/24 • Next hop st0.1
Security policy	through-vpn	Permit the following traffic: • From zone: ls-product-design-trust • To zone: ls-product-design-untrust • Source address: any • Destination address: 192.168.168.0/24 • Application: any

Configuration

IN THIS SECTION

- Procedure | 246

Procedure

CLI Quick Configuration

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

```
set interfaces st0 unit 1 family inet address 10.11.11.150/24
set routing-options static route 192.168.168.0/24 next-hop st0.1
set security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust policy
through-vpn match source-address any
```

```
set security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust policy
through-vpn match destination-address 192.168.168.0/24
set security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust policy
through-vpn match application any
set security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust policy
through-vpn then permit
```

Step-by-Step Procedure

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

To configure a route-based VPN tunnel in a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

```
[edit]
lsdesignadmin1@host:ls-product-design>configure
lsdesignadmin1@host:ls-product-design#
```

2. Configure the VPN tunnel interface.

```
[edit interfaces]
lsdesignadmin1@host:ls-product-design# set st0 unit 1 family inet address 10.11.11.150/24
```

3. Create a static route to the remote destination.

```
[edit routing-options]
lsdesignadmin1@host:ls-product-design# set static route 192.168.168.0/24 next-hop st0.1
```

4. Configure a security policy to permit traffic to the remote destination.

```
[edit security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust]
lsdesignadmin1@host:ls-product-design# set policy through-vpn match source-address any
lsdesignadmin1@host:ls-product-design# set policy through-vpn match destination-address
192.168.168.0/24
```

```
lsdesignadmin1@host:ls-product-design# set policy through-vpn match application any
lsdesignadmin1@host:ls-product-design# set policy through-vpn then permit
```

Results

From configuration mode, confirm your configuration by entering the `show interfaces st0`, `show routing-options`, and `show security policies` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
lsdesignadmin1@host:ls-product-design# show interfaces st0
    unit 1 {
        family inet {
            address 10.11.11.150/24;
        }
    }
    lsdesignadmin1@host:ls-product-design# show routing-options
        static {
            route 192.168.168.0/24 next-hop st0.1;
        }
    [edit]
    lsdesignadmin1@host:ls-product-design# show security policies
        from-zone ls-product-design-trust to-zone ls-product-design-untrust {
            policy through-vpn {
                match {
                    source-address any;
                    destination-address 192.168.168.0/24;
                    application any;
                }
                then {
                    permit;
                }
            }
            ...
        }
```

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

Verification

IN THIS SECTION

- Verifying the IKE Phase 1 Status | 249
- Verifying the IPsec Phase 2 Status | 249

Confirm that the configuration is working properly.

NOTE: Before starting the verification process, you need to send traffic from a host in the user logical system to a host in the 192.168.168.0/24 network. For example, initiate a ping from a host in the 12.1.1.0/24 subnet in the ls-product-design user logical system to the host 192.168.168.10.

Verifying the IKE Phase 1 Status

Purpose

Verify the IKE Phase 1 status.

Action

From operational mode, enter the `show security ike security-associations` command. After obtaining an index number from the command, use the `show security ike security-associations index` *`index_number`* `detail` command.

For sample outputs and meanings, see the “Verification” section of *Example: Configuring a Route-Based VPN*.

Verifying the IPsec Phase 2 Status

Purpose

Verify the IPsec Phase 2 status.

Action

From operational mode, enter the `show security ipsec security-associations` command. After obtaining an index number from the command, use the `show security ipsec security-associations index index_number detail` command.

For sample outputs and meanings, see the "Verification" section of *Example: Configuring a Route-Based VPN*.

SEE ALSO

- Example: Configuring a Route-Based VPN
- Understanding Route-Based VPN Tunnels in Logical Systems | **233**
- User Logical Systems Configuration Overview | **51**

RELATED DOCUMENTATION

- IPv6 Addresses in Logical Systems Overview | **360**

UTM for Logical Systems

IN THIS SECTION

- Understanding UTM Features in Logical Systems | **251**
- Example: Configuring UTM for the Primary Logical System | **252**
- Example: Configuring UTM for a User Logical System | **262**

Unified threat management (UTM) provides multiple security features and services for SRX Series devices on the network, protecting users from security threats in a simplified way. UTM secures the logical systems from viruses, malware, or malicious attachments by scanning the incoming data using Deep Packet Inspection and prevents access to unwanted websites by installing Enhanced Web Filtering (EWF).

Understanding UTM Features in Logical Systems

Unified Threat Management (UTM) in logical systems provides several security features such as antispam, antivirus, content filtering, and Web filtering to secure users from multiple Internet-borne threats. The advantage of UTM is streamlined installation and management of these multiple security capabilities. In logical systems the primary administrator configures the UTM features for the primary logical system. Configuring UTM features for logical systems is similar to configuring UTM features on a device that is not configured for logical systems.

The security features provided as part of the UTM solution are:

- *Antispam Filtering*—E-mail spam consists of unwanted e-mail messages, usually sent by commercial, malicious, or fraudulent entities. The antispam feature examines transmitted e-mail messages to identify e-mail spam. The default antispam feature is configured at the primary logical system and it is applicable for all the user logical systems.
- *Content Filtering*—Content filtering blocks or permits certain types of traffic based on the MIME type, file extension, protocol command, and embedded object type. The default content filtering feature is configured at the primary logical system and it is applicable for all the user logical systems.
- *Web Filtering*—Web filtering lets you manage Internet usage by preventing access to inappropriate Web content. The default Web filtering feature is configured at the primary logical system, and the user logical systems inherit these default Web filtering configuration.
- *Sophos Antivirus*—Sophos Antivirus scanning is offered as a less CPU-intensive alternative to the full file-based antivirus feature. Sophos Antivirus is as an in-the-cloud antivirus solution. The default antivirus feature is configured at the primary logical system, and the user logical systems inherit these default antivirus configuration.

You must configure the custom objects for the Web filtering, anti-spam, and content filtering features before configuring the UTM features. You can configure custom objects for each user logical system.

The predefined UTM default policy parameters for Web filtering, content filtering, antivirus, and antispam profiles are configured at the primary logical system. The user logical systems inherit the same antivirus and Web filtering features configured for the primary logical system. The options such as `mime-whitelist` and `url-whitelist` in antivirus profile, and `address-blacklist` and `address-whitelist` in antispam profile can be configured at the following hierarchy levels, respectively:

- `[edit security utm feature-profile anti-virus sophos-engine profile]`
- `[edit security utm feature-profile anti-spam sbl profile]`

The options `url-whitelist` and `url-blacklist` are not supported in the Web fiterling profile, you can use the custom category option to achieve the function.

Example: Configuring UTM for the Primary Logical System

IN THIS SECTION

- Requirements | 252
- Overview | 252
- Configuration | 253
- Verification | 258

This example shows how to configure the UTM features antivirus, antispam, content filtering, and Web filtering in the primary logical system. The primary administrator is responsible for assigning the UTM features to the user logical systems.

Requirements

This example uses the following hardware and software components:

- SRX Series device configured with the logical systems.
- Junos OS Release 18.3R1 and later releases.

Before you begin:

- Understand how to log in to the primary logical system as the primary administrator. See "Primary Logical Systems Overview" on page 23.
- Configure the interfaces, routing instances, and static routes for the primary logical system. See "Example: Configuring Interfaces, Routing Instances, and Static Routes for the Primary and Interconnect Logical Systems and Logical Tunnel Interfaces for the User Logical Systems (Primary Administrators Only)" on page 114.

Overview

By default, all system resources are assigned to the primary logical system, and the primary administrator allocates them to the user logical systems. The primary administrator manages the device and the logical systems.

This example shows how to configure the UTM features described in Table 21 on page 253 for the primary logical system.

Table 21: UTM Configuration Type, Steps, and Parameters

Configuration Type	Configuration Description	Configuration Parameter
Custom objects	Configure the MIME (Multipurpose Internet Mail Extension) types (my_blockmime01) to decide which traffic is allowed to bypass various types of scanning	*[multipart/ application/]*
	Define a set of file extensions (my_fileextlist01) that are used in file extension scan mode (scan-by-extension)	*[txt pl com zip]*
	Configure a URL pattern list (black_list) of URLs or addresses that you want to block.	*www.example.com*
	Configure a custom URL category (cust_black) of URLs or addresses that you want to block.	*black_list*
Antispam	Configure the antispam type server-based spam block list (SBL).	*sbl*
Antivirus	Configure the antivirus type Sophos Antivirus (sophos-engine) profile (mysav) scan option to scan specific types of traffic.	*uri-check*
Web filtering	Specify an action for Enhanced Web Filtering (EWF) (juniper-enhanced) profile (myewf), for requests that experience internal errors in the Web filtering module.	*log-and-permit*

In this procedure, you define custom objects, configure feature profiles for UTM features (antispam, antivirus, content filtering, and Web filtering), configure a UTM policy and attach feature profiles, and apply the UTM policy to the security policy as an application service. For more information, see the Unified Threat Management User Guide.

Configuration

IN THIS SECTION

CLI Quick Configuration | 254

Procedure | 255
Results | 256

CLI Quick Configuration

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

```
set security utm custom-objects mime-pattern my_blockmime01 value [ multipart/ application/ ]
set security utm custom-objects filename-extension my_fileextlist01 value [ txt pl com zip ]
set security utm custom-objects url-pattern black_list value www.example.com
set security utm custom-objects custom-url-category cust_black value black_list
set security utm default-configuration anti-virus type sophos-engine
set security utm default-configuration web-filtering type juniper-enhanced
set security utm default-configuration web-filtering juniper-enhanced cache timeout 1800
set security utm default-configuration web-filtering juniper-enhanced cache size 0
set security utm default-configuration anti-spam type sbl
set security utm feature-profile anti-virus sophos-engine profile mysav scan-options uri-check
set security utm feature-profile web-filtering juniper-enhanced profile myewf default log-and-
permit
set security utm utm-policy utm-p1 anti-virus http-profile mysav
set security utm utm-policy utm-p1 content-filtering http-profile junos-cf-defaults
set security utm utm-policy utm-p1 web-filtering http-profile myewf
set security utm utm-policy utm-p1 anti-spam smtp-profile junos-as-defaults
set security policies from-zone trust to-zone untrust policy p1 match source-address any
set security policies from-zone trust to-zone untrust policy p1 match destination-address any
set security policies from-zone trust to-zone untrust policy p1 match application junos-http
set security policies from-zone trust to-zone untrust policy p1 then permit application-services
utm-policy utm-p1
```

Procedure

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the Junos OS CLI User Guide.

1. Log in to the primary logical system as the primary administrator and enter configuration mode.

```
admin@host> configure
admin@host#
```

2. Configure the custom objects for the primary logical system.

```
[edit security utm custom-objects]
admin@host# set mime-pattern my_blockmime01 value [ multipart/ application/ ]
admin@host# set filename-extension my_fileextlist01 value [ txt pl com zip ]
admin@host# set url-pattern black_list value www.example.com
admin@host# set custom-url-category cust_black value black_list
```

3. Define the UTM default configuration for the primary logical system.

```
[edit security utm default-configuration]
admin@host# set anti-virus type sophos-engine
admin@host# set web-filtering type juniper-enhanced
admin@host# set web-filtering juniper-enhanced cache timeout 1800
admin@host# set web-filtering juniper-enhanced cache size 0
admin@host# set anti-spam type sbl
```

4. Configure the feature profile for the primary logical system.

```
[edit security utm feature-profile]
admin@host# set anti-virus sophos-engine profile mysav scan-options uri-check
admin@host# set web-filtering juniper-enhanced profile myewf default log-and-permit
```

5. Configure the UTM policy for the primary logical system.

```
[edit security utm utm-policy]
admin@host# set utm-p1 anti-virus http-profile mysav
admin@host# set utm-p1 content-filtering http-profile junos-cf-defaults
admin@host# set utm-p1 web-filtering http-profile myewf
admin@host# set utm-p1 anti-spam smtp-profile junos-as-defaults
```

6. Configure the security policies for the primary logical system.

```
[edit security policies]
admin@host# set from-zone trust to-zone untrust policy p1 match source-address any
admin@host# set from-zone trust to-zone untrust policy p1 match destination-address any
admin@host# set from-zone trust to-zone untrust policy p1 match application junos-http
admin@host# set from-zone trust to-zone untrust policy p1 permit application-services utm-
policy utm-p1
```

Results

From configuration mode, confirm your configuration by entering the `show security` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
admin@host# show security
utm {
    custom-objects {
        mime-pattern {
            my_blockmime01 {
                value [ multipart/ application/ ];
            }
        }
        filename-extension {
            my_fileextlist01 {
                value [ txt pl com zip ];
            }
        }
        url-pattern {
            black_list {
                value www.example.com;
            }
```

```
        }
        custom-url-category {
            cust_black {
                value black_list;
            }
        }
    }
    default-configuration {
        anti-virus {
            type sophos-engine;
        }
        web-filtering {
            type juniper-enhanced;
            juniper-enhanced {
                cache {
                    timeout 1800;
                    size 0;
                }
            }
        }
        anti-spam {
            type sbl;
        }
    }
    feature-profile {
        anti-virus {
            sophos-engine {
                profile mysav {
                    scan-options {
                        uri-check;
                    }
                }
            }
        }
        web-filtering {
            juniper-enhanced {
                profile myewf {
                    default log-and-permit;
                }
            }
        }
    }
    utm-policy utm-p1 {
```

```
        anti-virus {
            http-profile mysav;
        }
        content-filtering {
            http-profile junos-cf-defaults;
        }
        web-filtering {
            http-profile myewf;
        }
        anti-spam {
            smtp-profile junos-as-defaults;
        }
    }
}
policies {
    from-zone trust to-zone untrust {
        policy p1 {
            match {
                source-address any;
                destination-address any;
                application junos-http;
            }
            then {
                permit {
                    application-services {
                        utm-policy utm-p1;
                    }
                }
            }
        }
    }
}
```

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

Verification

IN THIS SECTION

- Verifying Antivirus Configuration | **259**
- Verifying Antispam Configuration | **260**

Verifying Content Filtering Configuration | 261
Verifying Web Filtering Configuration | 261

To confirm that the configuration is working properly, perform these tasks:

Verifying Antivirus Configuration

Purpose

Verify that the antivirus feature is configured for the primary logical system.

Action

From operational mode, enter the `show security utm anti-virus statistics` command to view the details of the antivirus feature configured for the primary logical system.

```
admin@host> show security utm anti-virus statistics
 UTM Anti Virus statistics:
 MIME-whitelist passed:                0
URL-whitelist passed:                 0
Session abort:                        0
Scan Request:

  Total           Clean         Threat-found    Fallback
      9               7              1               1

Fallback:
                       Log-and-Permit    Block       Permit
Engine not ready:          0                0           0
 Out of resources:         0                0           0
Timeout:                   0                0           0
Maximum content size:      1                0           0
Too many requests:         0                0           0
Others:                    0                0           0
```

Meaning

The output displays the antisvirus statistics for the primary logical system.

Verifying Antispam Configuration

Purpose

Verify that the antispam feature is configured for the primary logical system.

Action

From operational mode, enter the `show security utm anti-spam statistics` command to view the details of the antispam feature configured for the primary logical system.

```
admin@host> show security utm anti-spam statistics
 UTM Anti Spam statistics:

Total connections:     1
Denied connections:    1
Total greetings:       0
Denied greetings:      0
Total e-mail scanned:  0
White list hit:        0
Black list hit:        0
Spam total:            0
Spam tagged:           0
Spam dropped:          0
DNS errors:            0
Timeout errors:        0
Return errors:         0
Invalid parameter errors: 0
```

Meaning

The output displays the antispam statistics for the primary logical system.

Verifying Content Filtering Configuration

Purpose

Verify that the content filtering feature is configured for the primary logical system.

Action

From operational mode, enter the `show security utm content-filtering statistics` command to view the details of the content filtering feature configured for the primary logical system.

```
admin@host> show security content-filtering statistics
 Content-filtering-statistic:         Blocked
     Base on command list:                    0
     Base on mime list:                       1
     Base on extension list:                  0
     ActiveX plugin:                          0
     Java applet:                             0
     EXE files:                               0
     ZIP files:                               0
     HTTP cookie:                             0
```

Meaning

The output displays the content filtering statistics for the primary logical system.

Verifying Web Filtering Configuration

Purpose

Verify that the Web filtering feature is configured for the primary logical system.

Action

From operational mode, enter the `show security utm web-filtering statistics` command to view the details of the Web filtering feature configured for the primary logical system.

```
admin@host> show security web-filtering statistics
 UTM web-filtering statistics:
    Total requests:                    4
```

```
    white list hit:                     1
    Black list hit:                     1
    Custom category permit:             1
    Custom category block:              1
    Custom category quarantine:         0
    Custom category qurantine block:    0
    Custom category quarantine permit:  0
    Web-filtering sessions in total:    64000
    Web-filtering sessions in use:      0
    Fallback:                   log-and-permit     block
          Default                             0         0
          Timeout                             0         0
     Connectivity                             0         0
Too-many-requests                             0         0
```

Meaning

The output displays the Web filtering statistics for the primary logical system.

Example: Configuring UTM for a User Logical System

IN THIS SECTION

- Requirements | 262
- Overview | 263
- Configuration | 264
- Verification | 269

This example shows how to configure the UTM features antivirus, antispam, content filtering, and Web filtering for a user logical system. The primary administrator creates a user logical system and assigns an administrator for managing the user logical system. A user logical system can have multiple user logical system administrators.

Requirements

This example uses the following hardware and software components:

- SRX Series device configured with the logical systems.
- Junos OS Release 18.3R1 and later releases.

Before you begin:

- Understand the user logical system administrator role and functions. See "Understanding User Logical Systems and the User Logical System Administrator Role" on page 53.
- Understand how to log in to the user logical system as an administrator. See "User Logical Systems Configuration Overview" on page 51.
- This example shows how to configure the UTM features for the ls-product-design user logical system. To understand how to create the ls-product-design user logical system, see "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System" on page 57.

Overview

The primary administrator assigns the UTM features antivirus, antispam, content filtering, and Web filtering to the user logical system. The user logical system administrator can configure and manage the UTM features for the user logical systems. The antispam, antivirus and Web filtering features are configured in the primary logical system are described in Table 22 on page 263. All the user logical systems can use the same antispam, antivirus and Web filtering features with the same profile.

Table 22: UTM Configuration Type, Steps, and Parameters

Configuration Type	Configuration Description	Configuration Parameter
Custom objects	Configure a URL pattern (`url1`) of URL patterns that bypass scanning.	*www.abc.com*
	Configure a custom URL category (`cust1`) of URLs or addresses list that bypass scanning.	*url1*
	Configure a custom message type (`redirect-url`) to redirect traffic destined for protected sources.	*http://www.example1.com.cn*
Antispam	Configure antispam profile (`as1`) spam action.	*block*

Table 22: UTM Configuration Type, Steps, and Parameters *(Continued)*

Configuration Type	Configuration Description	Configuration Parameter
Antivirus	Configure antivirus profile (sav1) fallback option.	*log-and-permit*
	Configure antivirus profile (sav1) scan option.	*uri-check*
Web filtering	Configure Web filtering profile (ewf1) category (cust1) action.	*block*
	Configure Web filtering profile (ewf1) category (cust1) custom message.	*custmsg1*
	Configure Web filtering profile (ewf1) category (Enhanced_Search_Engines_and_Portals) action.	*block*
	Specify an action for Enhanced Web Filtering (EWF) (juniper-enhanced) profile (ewf1), for requests that experience internal errors in the Web filtering module.	*log-and-permit*

Configuration

IN THIS SECTION

- CLI Quick Configuration | **265**
- Procedure | **266**
- Results | **267**

CLI Quick Configuration

To quickly configure this example, log in to the ls-product-design user logical system as the administrator, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the `[edit]` hierarchy level, and then enter `commit` from configuration mode.

```
set security utm custom-objects url-pattern url1 value www.abc.com
set security utm custom-objects custom-url-category cust1 value url1
set security utm custom-objects custom-message cust-msg1 type redirect-url content http://
www.example1.com.cn
set security utm feature-profile anti-virus sophos-engine profile sav1 fallback-options default
log-and-permit
set security utm feature-profile anti-virus sophos-engine profile sav1 scan-options uri-check
set security utm feature-profile web-filtering juniper-enhanced profile ewf1 category cust1
action block
set security utm feature-profile web-filtering juniper-enhanced profile ewf1 category cust1
custom-message custmsg1
set security utm feature-profile web-filtering juniper-enhanced profile ewf1 category
Enhanced_Search_Engines_and_Portals action block
set security utm feature-profile web-filtering juniper-enhanced profile ewf1 default log-and-
permit
set security utm feature-profile web-filtering juniper-enhanced profile ewf2 default log-and-
permit
set security utm feature-profile anti-spam sbl profile as1 spam-action block
set security utm utm-policy utm-p1 anti-virus http-profile sav1
set security utm utm-policy utm-p1 web-filtering http-profile juniper-enhanced
set security utm utm-policy utm-p1 anti-spam smtp-profile as1
set security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust policy
sec_policy match source-address any
set security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust policy
sec_policy match destination-address any
set security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust policy
sec_policy match application any
set security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust policy
sec_policy then permit application-services utm-policy utm-p1
```

Procedure

Step-by-Step Procedure

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

1. Log in to the ls-product-design user logical system as the administrator and enter configuration mode.

```
lsdesignadmin1@host:ls-product-design> configure
lsdesignadmin1@host:ls-product-design#
```

2. Configure the custom objects for the ls-product-design user logical system.

```
[edit security utm custom-objects]
lsdesignadmin1@host:ls-product-design# set url-pattern url1 value www.abc.com
lsdesignadmin1@host:ls-product-design# set custom-url-category cust1 value url1
lsdesignadmin1@host:ls-product-design# set custom-message cust-msg1 type redirect-url content
http://www.example1.com.cn
```

3. Configure the feature profiles for the ls-product-design user logical system.

```
[edit security utm feature-profile]
lsdesignadmin1@host:ls-product-design# set anti-virus sophos-engine profile sav1 fallback-
options default log-and-permit
lsdesignadmin1@host:ls-product-design# set anti-virus sophos-engine profile sav1 scan-options
uri-check
lsdesignadmin1@host:ls-product-design# set web-filtering juniper-enhanced profile ewf1
category cust1 action block
lsdesignadmin1@host:ls-product-design# set web-filtering juniper-enhanced profile ewf1
category cust1 custom-message custmsg1
lsdesignadmin1@host:ls-product-design# set web-filtering juniper-enhanced profile ewf1
category Enhanced_Search_Engines_and_Portals action block
lsdesignadmin1@host:ls-product-design# set web-filtering juniper-enhanced profile ewf1
default log-and-permit
lsdesignadmin1@host:ls-product-design# set web-filtering juniper-enhanced profile ewf2
default log-and-permit
lsdesignadmin1@host:ls-product-design# set anti-spam sbl profile as1 spam-action block
```

4. Configure the UTM policy for the ls-product-design user logical system.

```
[edit security utm utm-policy]
lsdesignadmin1@host:ls-product-design# set utm-p1 anti-virus http-profile sav1
lsdesignadmin1@host:ls-product-design# set utm-p1 web-filtering http-profile juniper-enhanced
lsdesignadmin1@host:ls-product-design# set utm-p1 anti-spam smtp-profile as1
```

5. Configure the security policies for the ls-product-design user logical system.

```
[edit security policies]
lsdesignadmin1@host:ls-product-design# set from-zone lsys1-trust to-zone lsys1-untrust policy
sec_policy match source-address any
lsdesignadmin1@host:ls-product-design# set from-zone lsys1-trust to-zone lsys1-untrust policy
sec_policy match destination-address any
lsdesignadmin1@host:ls-product-design# set from-zone lsys1-trust to-zone lsys1-untrust policy
sec_policy match application any
lsdesignadmin1@host:ls-product-design# set from-zone lsys1-trust to-zone lsys1-untrust policy
sec_policy then permit application-services utm-policy utm-p1
```

Results

From configuration mode, confirm your configuration by entering the `show security` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
lsdesignadmin1@host:ls-product-design# show security
utm {
    custom-objects {
        url-pattern {
            url1 {
                value www.abc.com;
            }
        }
        custom-url-category {
            cust1 {
                value url1;
            }
        }
        custom-message {
            cust-msg1 {
```

```
                type redirect-url;
                content http://www.example1.com.cn;
            }
        }
    }
    feature-profile {
        anti-virus {
            sophos-engine {
                profile sav1 {
                    fallback-options {
                        default log-and-permit;
                    }
                    scan-options {
                        uri-check;
                    }
                }
            }
        }
        web-filtering {
            juniper-enhanced {
                profile ewf1 {
                    category {
                        cust1 {
                            action block;
                            custom-message custmsg1;
                        }
                        Enhanced_Search_Engines_and_Portals {
                            action block;
                        }
                    }
                    default log-and-permit;
                }
                profile ewf2 {
                    default log-and-permit;
                }
            }
        }
        anti-spam {
            sbl {
                profile as1 {
                    spam-action block;
                }
            }
```

```
        }
    }
    utm-policy utm-p1 {
        anti-virus {
            http-profile sav1;
        }
        web-filtering {
            http-profile juniper-enhanced;
        }
        anti-spam {
            smtp-profile as1;
        }
    }
}
policies {
    from-zone ls-product-design-trust to-zone ls-product-design-untrust {
        policy sec_policy {
            match {
                source-address any;
                destination-address any;
                application any;
            }
            then {
                permit {
                    application-services {
                        utm-policy utm-p1;
                    }
                }
            }
        }
    }
}
```

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

Verification

IN THIS SECTION

- Verifying Antivirus Configuration | **270**
- Verifying Antispam Configuration | **271**

Verifying Content Filtering Configuration | 272
Verifying Web Filtering Configuration | 272

To confirm that the configuration is working properly, perform these tasks:

Verifying Antivirus Configuration

Purpose

Verify that the antivirus feature is configured for the ls-product-design user logical system.

Action

From operational mode, enter the `show security utm anti-virus statistics` command to view the antivirus statistics information for the ls-product-design user logical system.

```
lsdesignadmin1@host:ls-product-design> show security utm anti-virus statistics
 UTM Anti Virus statistics:
 MIME-whitelist passed:                0
URL-whitelist passed:                 0
Session abort:                        0
Scan Request:

  Total           Clean         Threat-found    Fallback
    9               7             1               1

Fallback:
                       Log-and-Permit    Block       Permit
Engine not ready:          0                0           0
 Out of resources:         0                0           0
Timeout:                   0                0           0
Maximum content size:      1                0           0
Too many requests:         0                0           0
Others:
```

Meaning

The output displays the antisvirus statistics information for the ls-product-design user logical system.

Verifying Antispam Configuration

Purpose

Verify that the antispam feature is configured for the ls-product-design user logical system.

Action

From operational mode, enter the `show security utm anti-spam statistics` command to view the antispam statistics information for the ls-product-design user logical system.

```
lsdesignadmin1@host:ls-product-design> show security utmnti-spam statistics
 UTM Anti Spam statistics:

Total connections:     1
Denied connections:    1
Total greetings:       0
Denied greetings:      0
Total e-mail scanned:  0
White list hit:        0
Black list hit:        0
Spam total:            0
Spam tagged:           0
Spam dropped:          0
DNS errors:            0
Timeout errors:        0
Return errors:         0
Invalid parameter errors: 0
```

Meaning

The output displays the antispam statistics information for the ls-product-design user logical system.

Verifying Content Filtering Configuration

Purpose

Verify that the content filtering feature is configured for the ls-product-design user logical system.

Action

From operational mode, enter the `show security utm content-filtering statistics` command to view the content filtering statistics information for the ls-product-design user logical system.

```
lsdesignadmin1@host:ls-product-design> show security content-filteringtatistics
 Content-filtering-statistic:         Blocked
     Base on command list:                    0
     Base on mime list:                       1
     Base on extension list:                  0
     ActiveX plugin:                          0
     Java applet:                             0
     EXE files:                               0
     ZIP files:                               0
     HTTP cookie:                             0
```

Meaning

The output displays the content filtering statistics information for the ls-product-design user logical system.

Verifying Web Filtering Configuration

Purpose

Verify that the Web filtering feature is configured for the ls-product-design user logical system.

Action

From operational mode, enter the `show security utm web-filtering statistics` command to view the Web filtering statistics information for the ls-product-design user logical system.

```
lsdesignadmin1@host:ls-product-design> show security web-filteringstatistics
 UTM web-filtering statistics:
```

```
    Total requests:                    4
    white list hit:                    1
    Black list hit:                    1
    Custom category permit:            1
    Custom category block:             1
    Custom category quarantine:        0
    Custom category qurantine block:   0
    Custom category quarantine permit: 0
    Web-filtering sessions in total:   64000
    Web-filtering sessions in use:     0
    Fallback:         log-and-permit          block
          Default                  0              0
          Timeout                  0              0
     Connectivity                  0              0
Too-many-requests                  0              0
```

Meaning

The output displays the Web filtering statistics information for the ls-product-design user logical system.

IDP for Logical Systems

IN THIS SECTION

- IDP in Logical Systems Overview | **274**
- Understanding IDP Features in Logical Systems | **276**
- Example: Configuring an IDP Policy for the Primary Logical Systems | **280**
- Example: Configuring and Assigning a Predefined IDP Policy for a User Logical System | **288**
- Example: Enabling IDP in a User Logical System Security Policy | **292**
- Example: Configuring an IDP Policy for a User Logical System | **296**

An Intrusion Detection and Prevention (IDP) policy in logical systems enables you to selectively enforce various attack detection and prevention techniques on the network traffic passing through your SRX Series. The SRX Series offer the same set of IDP signatures that are available on Juniper Networks IDP

Series Intrusion Detection and Prevention Appliances to secure networks against attacks. For more information, see the following topics:

IDP in Logical Systems Overview

IN THIS SECTION

- IDP Policies | 274
- Limitation | 276
- IDP Installation and Licensing for Logical Systems | 276

A Junos OS Intrusion Detection and Prevention (IDP) policy enables you to selectively enforce various attack detection and prevention techniques on network traffic passing through a logical system.

This topic includes the following sections:

IDP Policies

The primary administrator configures IDP policies at the root level. Configuring an IDP policy for logical systems is similar to configuring an IDP policy on a device that is not configured for logical systems. This can include the configuration of custom attack objects.

IDP policy templates installed in root logical system are visible and used by all logical systems.

The primary administrator then specifies an IDP policy in the security profile that is bound to a logical system. To enable IDP in a logical system, the primary administrator or user logical system administrator configures a security policy that defines the traffic to be inspected and specifies the `permit application-services idp` action.

Although the primary administrator can configure multiple IDP policies, a logical system can have only one active IDP policy at a time. For user logical systems, the primary administrator can either bind the same IDP policy to multiple user logical systems or bind a unique IDP policy to each user logical system. To specify the active IDP policy for the primary logical system, the primary administrator can *either* reference the IDP policy in the security profile that is bound to the primary logical system or use the `active-policy` *configuration statement* at the [`edit security idp`] hierarchy level.

The root administrator configures the number of maximum IDP sessions reservation for a root and user logical system. The number of IDP sessions that are allowed for a root logical system are defined using

the command `set security idp max-sessions` *`max-sessions`* and the number of IDP sessions that are allowed for a user logical system are defined using the command `set security idp logical-system` *`logical-system`* `max-sessions` *`max-sessions`*.

NOTE: A commit error is generated if an IDP policy is both configured in the security profile that is bound to the primary logical system and specified with the `active-policy` configuration statement. Use only one method to specify the active IDP policy for the primary logical system.

NOTE: If you have configured more than one IDP policy in a security policy, then configuring default IDP policy configuration is mandatory.

A default IDP policy configuration is supported when multiple IDP policies are available. The default IDP policy is one of the multiple IDP policies. For more information about configuring multiple IDP policies and default IDP policy, see the *IDP Policy Selection for Unified Policies*.

The logical system administrator performs the following actions:

- Configure multiple IDP policies and attach to the firewall policies to be used by the user logical systems. If the IDP policy is not configured for a user logical system, the default IDP policy configured by the primary administrator is used. The IDP policy is bound to the user logical systems through a logical systems security policy.
- Create or modify IDP policies for their user logical systems. The IDP policies are bound to user logical systems. When an IDP policy is changed, and commit succeeds, the existing sessions mapped to current active policy continue to use the old IDP combined policy. When an IDP policy is changed, and commit fails, only the logical system user that has initiated the commit change is notified about the commit failure.
- The logical system can create security zones in the user logical system and assign interfaces to each security zone. Zones that are specific to user logical systems cannot be referenced in IDP policies configured by the primary administrator. The primary administrator can reference zones in the primary logical system in an IDP policy configured for the primary logical system.
- View the attack statistics detected and IDP counters, attack table, and policy commit status by the individual logical system using the commands `show security idp counters`, `show security idp attack table`, `show security idp policies`, `show security idp policy-commit-status`, and `show security idp security-package-version`.

Limitation

- When a IDP policy is changed and compiled in a specific user logical system, this change is considered as a single global policy change and compiled for all policies of all the logical systems.

IDP Installation and Licensing for Logical Systems

An idp-sig license must be installed at the root level. Once IDP is enabled at the root level, it can be used with any logical system on the device.

A single IDP security package is installed for all logical systems on the device at the root level. The download and install options can only be executed at the root level. The same version of the IDP attack database is shared by all logical systems.

SEE ALSO

User Logical Systems Configuration Overview | 51

Understanding Logical Systems Security Profiles (Primary Administrators Only) | 71

IDP Policies Overview

Understanding IDP Features in Logical Systems

IN THIS SECTION

- Rulebases | 277
- Protocol Decoders | 277
- SSL Inspection | 277
- Inline Tap Mode | 277
- Multi-Detectors | 278
- Logging and Monitoring | 278

This topic includes the following sections:

Rulebases

A single IDP policy can contain only one instance of any type of rulebase. The following IDP rulebases are supported for logical systems:

- The Intrusion prevention system (IPS) rulebase uses attack objects to detect known and unknown attacks. It detects attacks based on stateful signature and protocol anomalies.
- The application-level distributed denial-of-service (DDoS) rulebase defines parameters to protect servers such as DNS or HTTP. The application-level DDoS rulebase defines the source match condition for traffic that should be monitored and takes an action, such as drop the connection, drop the packet, or no action. It can also perform actions against future connections that use the same IP address.

NOTE: Status monitoring for IPS and application-level DDoS is global to the device and not on a per logical system basis.

Protocol Decoders

The Junos IDP module ships with a set of preconfigured protocol decoders. These protocol decoders have default settings for various protocol-specific contextual checks that they perform. The IDP protocol decoder configuration is global and applies to all logical systems. Only the primary administrator at the root level can modify the settings at the [`edit security idp sensor-configuration`] hierarchy level.

SSL Inspection

IDP SSL inspection uses the Secure Sockets Layer (SSL) protocol suite to enable inspection of HTTP traffic encrypted in SSL.

SSL inspection configuration is global and applies to all logical systems on a device. SSL inspection can only be configured by the primary administrator at the root level with the `ssl-inspection` *configuration statement* at the [`edit security idp sensor-configuration`] hierarchy level.

Inline Tap Mode

The inline tap mode feature provides passive, inline detection of Application Layer threats for traffic matching security policies that have the IDP application service enabled. When a device is in inline tap mode, packets pass through firewall inspection and are also copied to the independent IDP module. This allows the packets to get to the next service module without waiting for IDP processing results.

Inline tap mode is enabled or disabled for all logical systems at the root level by the primary administrator. To enable inline tap mode, use the `inline-tap` configuration statement at the [`edit security forwarding-process application-services maximize-idp-sessions`] hierarchy level. Delete the inline tap mode configuration to switch the device back to regular mode.

NOTE: The device must be restarted when switching to inline tap mode or back to regular mode.

Multi-Detectors

When a new IDP security package is received, it contains attack definitions and a detector. After a new policy is loaded, it is also associated with a detector. If the policy being loaded has an associated detector that matches the detector already in use by the existing policy, the new detector is not loaded and both policies use a single associated detector. But if the new detector does not match the current detector, the new detector is loaded along with the new policy. In this case, each loaded policy will then use its own associated detector for attack detection.

The version of the detector is common to all logical systems.

Logging and Monitoring

Status monitoring options are available to the primary administrator only. All status monitoring options under the `show security idp` and `clear security idp` CLI operational commands present global information, but not on a per logical system basis.

NOTE: SNMP monitoring for IDP is not supported on logical systems.

IDP generates event logs when an event matches an IDP policy rule in which logging is enabled.

The logical systems identification is added to the following types of IDP traffic processing logs:

- Attack logs. The following example shows an attack log for the ls-product-design logical system:

```
Feb 22 14:06:00 aqgpo1ifw01 RT_IDP: %-IDP_ATTACK_LOG_EVENT_LS: Lsys A01: IDP: At 1329883555,
ANOMALY Attack log <10.1.128.200/33699->192.168.22.84/80> for TCP protocol and service HTTP
application NONE by rule 4 of rulebase IPS in policy Policy1. attack: repeat=3, action=NONE,
threat-severity=INFO, name=HTTP:AUDIT:URL, NAT <0.0.0.0:0->0.0.0.0:0>, time-elapsed=0,
inbytes=0, outbytes=0, inpackets=0, outpackets=0, intf:NSS-Mgmt:reth0.55->SIEM-MGMT:reth0.60,
packet-log-id: 0 and misc-message
```

NOTE: In the IDP attack detection event log message (IDP_ATTACK_LOG_EVENT_LS), the time-elapsed, inbytes, outbytes, inpackets, and outpackets fields are not populated.

- IP action logs. The following example shows an IP action log for the ls-product-design logical system:

```
Oct 13 16:56:04 8.0.0.254 RT_IDP: IDP_ATTACK_LOG_EVENT_LS: IDP: In ls-product-design at
1287014163, TRAFFIC Attack log <25.0.0.1/34802->15.0.0.1/21> for TCP protocol and service
SERVICE_NONE application NONE by rule 1 of rulebase IPS in policy Recommended. attack:
repeat=0, action=TRAFFIC_IPACTION_NOTIFY, threat-severity=INFO, name=_, NAT <0.0.0.0:0-
>0.0.0.0:0>, time-elapsed=0, inbytes=0, outbytes=0, inpackets=0, outpackets=0, intf:ls-
product-design-trust:ge-0/0/1.0->ls-product-design-untrust:plt0.3, packet-log-id: 0 and misc-
message -
```

- Application DDoS logs. The following example shows an application DDoS log for the ls-product-design logical system:

```
Oct 11 16:29:57 8.0.0.254 RT_IDP: IDP_APPDDOS_APP_ATTACK_EVENT_LS: DDOS Attack in ls-product-
design at 1286839797 on my-http,
<ls-product-design-untrust:ge-0/0/0.0:4.0.0.1:33738->ls-product-design-
trust:ge-0/0/1.0:5.0.0.1:80> for TCP protocol and service HTTP by rule 1 of rulebase DDOS in
policy Recommended. attack: repeats 0 action DROP threat-severity INFO, connection-hit-rate
0,  context-name http-url-parsed, hit-rate 6, value-hit-rate 6 time-scope PEER time-count 2
time-period 10 secs, context value:  ascii: /abc.html hex: 2f 61 62 63 2e 68 74 6d 6c
```

SEE ALSO

Understanding IDP Policy Rule Bases

Understanding IDP Protocol Decoders

IDP SSL Overview

Understanding IDP Inline Tap Mode

Understanding Multiple IDP Detector Support

Understanding IDP Logging

Example: Configuring an IDP Policy for the Primary Logical Systems

IN THIS SECTION

- Requirements | 280
- Overview | 280
- Configuration | 282
- Verification | 287

This example shows how to configure an IDP policy in a primary logical system.

Requirements

Before you begin:

- Log in to the primary logical system as the primary administrator. See "Understanding the Primary Logical Systems and the Primary Administrator Role" on page 23.
- Read "IDP in Logical Systems Overview" on page 274.
- Use the `show system security-profile` command to see the resources allocated to the primary logical system.

Overview

In this example you configure a custom attack that is used in an IDP policy. The IDP policy is specified in a security profile that is applied to the primary logical system. IDP is then enabled in a security policy configured in the primary logical system.

You configure the features described in Table 23 on page 281.

Table 23: IDP Configuration for the Primary Logical System

Feature	Name	Configuration Parameters

Custom attack	http-bf	• Severity critical
		• Detect three attacks between source and destination addresses of sessions.
		• Stateful signature attack type with the following characteristics:
		  • location http-url-parsed
		  • pattern .*juniper.*
		  • client to server traffic

IPS rulebase policy	root-idp-policy	Match: • application default • http-bf custom attacks Action: • drop-connection • notification log-attacks
Logical system security profile	primary-profile (previously configured and applied to root-logical-system)	Add IDP policy root-idp-policy.
Security policy	enable-idp	Enable IDP in a security policy that matches any traffic from the lsys-root-untrust zone to the lsys-root-trust zone.

NOTE: A logical system can have only one active IDP policy at a time. To specify the active IDP policy for the primary logical system, the primary administrator can reference the IDP policy in the security profile that is bound to the primary logical system as shown in this example.

Alternatively, the primary administrator can use the `active-policy` configuration statement at the `[edit security idp]` hierarchy level.

A commit error is generated if an IDP policy is both configured in the security profile that is bound to the primary logical system and specified with the `active-policy` configuration statement. Use only one method to specify the active IDP policy for the primary logical system.

Configuration

IN THIS SECTION

- Configuring a Custom Attack | 282
- Configuring an IDP Policy for the Primary Logical System | 284
- Enabling IDP in a Security Policy | 286

Configuring a Custom Attack

CLI Quick Configuration

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

```
set security idp custom-attack http-bf severity critical
set security idp custom-attack http-bf time-binding count 3
set security idp custom-attack http-bf time-binding scope peer
set security idp custom-attack http-bf attack-type signature context http-url-parsed
set security idp custom-attack http-bf attack-type signature pattern .*juniper.*
set security idp custom-attack http-bf attack-type signature direction client-to-server
```

Step-by-Step Procedure

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

To configure a custom attack object:

1. Log in to the primary logical system as the primary administrator and enter configuration mode.

```
[edit]
admin@host> configure
admin@host#
```

2. Create the custom attack object and set the severity level.

```
[edit security idp]
admin@host# set custom-attack http-bf severity critical
```

3. Configure attack detection parameters.

```
[edit security idp]
admin@host# set custom-attack http-bf time-binding count 3
admin@host# set custom-attack http-bf time-binding scope peer
```

4. Configure stateful signature parameters.

```
[edit security idp]
admin@host# set custom-attack http-bf attack-type signature context http-url-parsed
admin@host# set custom-attack http-bf attack-type signature pattern .*juniper.*
admin@host# set custom-attack http-bf attack-type signature direction client-to-server
```

Results

From configuration mode, confirm your configuration by entering the `show security idp custom-attack http-bf` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
admin@host# show security idp custom-attack http-bf
severity critical;
    time-binding {
        count 3;
        scope peer;
    }
```

```
        attack-type {
            signature {
                context http-url-parsed;
                pattern .*juniper.*;
                direction client-to-server;
            }
        }
```

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

Configuring an IDP Policy for the Primary Logical System

CLI Quick Configuration

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

```
set security idp idp-policy root-idp-policy rulebase-ips rule 1 match application default
set security idp idp-policy root-idp-policy rulebase-ips rule 1 match attacks custom-attacks
http-bf
set security idp idp-policy root-idp-policy rulebase-ips rule 1 then action drop-connection
set security idp idp-policy root-idp-policy rulebase-ips rule 1 then notification log-attacks
set system security-profile master-profile idp-policy root-idp-policy
```

Step-by-Step Procedure

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

To configure an IDP policy:

1. Create the IDP policy and configure match conditions.

```
[edit security idp]
admin@host# set idp-policy root-idp-policy rulebase-ips rule 1 match application default
admin@host# set idp-policy root-idp-policy rulebase-ips rule 1 match attacks custom-attacks
http-bf
```

2. Configure actions for the IDP policy.

```
[edit security idp]
admin@host# set idp-policy root-idp-policy rulebase-ips rule 1 then action drop-connection
admin@host# set idp-policy root-idp-policy rulebase-ips rule 1 then notification log-attacks
```

3. Add the IDP policy to the security profile.

```
[edit system security-profile master-profile]
admin@host# set idp-policy lsys1-idp-policy
```

Results

From configuration mode, confirm your configuration by entering the `show security idp idp-policy root-idp-policy` and `show system security-profile master-profile` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
admin@host# show security idp idp-policy root-idp-policy
    rulebase-ips {
        rule 1 {
            match {
                application default;
                attacks {
                    custom-attacks http-bf;
                }
            }
            then {
                action {
                    drop-connection;
                }
                notification {
                    log-attacks;
                }
            }
        }
    }
admin@host# show system security-profile master-profile
```

```
...
idp-policy lsys1-idp-policy;
```

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

Enabling IDP in a Security Policy

CLI Quick Configuration

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

```
set security policies from-zone lsys-root-untrust to-zone lsys-root-trust policy enable-idp
match source-address any
set security policies from-zone lsys-root-untrust to-zone lsys-root-trust policy enable-idp
match destination-address any
set security policies from-zone lsys-root-untrust to-zone lsys-root-trust policy enable-idp
match application any
set security policies from-zone lsys-root-untrust to-zone lsys-root-trust policy enable-idp then
permit application-services idp
```

Step-by-Step Procedure

To enable IDP in a security policy:

1. Create the security policy and configure match conditions.

```
[edit security policies from-zone lsys-root-untrust to-zone lsys-root-trust]
admin@host# set policy enable-idp match source-address any
admin@host# set policy enable-idp match destination-address any
admin@host# set policy enable-idp match application any
```

2. Enable IDP.

```
[edit security policies from-zone lsys-root-untrust to-zone lsys-root-trust]
admin@host# set policy enable-idp then permit application-services idp
```

Results

From configuration mode, confirm your configuration by entering the `show security policies` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this `show` command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
    admin@host# show security policies
    from-zone lsys-root-untrust to-zone lsys-root-trust {
        policy enable-idp {
            match {
                source-address any;
                destination-address any;
                application any;
            }
            then {
                permit {
                    application-services {
                        idp;
                    }
                }
            }
        }
    }
...
```

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

Verification

IN THIS SECTION

- Verifying Attack Matches | 288

Verifying Attack Matches

Purpose

Verify that attacks are being matched in network traffic.

Action

From operational mode, enter the `show security idp attack table` command.

```
admin@host> show security idp attack table
IDP attack statistics:
  Attack name                                #Hits
  http-bf                                        1
```

SEE ALSO

IDP in Logical Systems Overview | 274

SRX Series Logical Systems Primary Administrator Configuration Tasks Overview | 24

Example: Configuring and Assigning a Predefined IDP Policy for a User Logical System

IN THIS SECTION

- Requirements | 289
- Overview | 289
- Configuration | 289
- Verification | 291

The primary administrator can *either* download predefined IDP policies to the device or configure custom IDP policies at the root level using custom or predefined attack objects. The primary

administrator is responsible for assigning an IDP policy to a user logical system. This example shows how to assign a predefined IDP policy to a user logical system.

Requirements

Before you begin:

- Log in to the primary logical system as the primary administrator. See "Understanding the Primary Logical Systems and the Primary Administrator Role" on page 23.
- Read *IDP Policies Overview*.
- Assign the ls-design-profile security policy to the ls-product-design user logical system. See "Example: Configuring Logical Systems Security Profiles (Primary Administrators Only)" on page 77.
- Download predefined IDP policy templates to the device. See *Downloading and Using Predefined IDP Policy Templates (CLI Procedure)*.

> **NOTE**: Activating a predefined IDP policy with the `active-policy` configuration statement at the [`edit security idp`] hierarchy level only applies to the primary logical system. For a user logical system, the primary administrator specifies the active IDP policy in the security profile that is bound to the user logical system.

Overview

The predefined IDP policy named Recommended contains attack objects recommended by Juniper Networks. All rules in the policy have their actions set to take the recommended action for each attack object. You add the Recommended IDP policy to the ls-design-profile, which is bound to the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System" on page 57.

Configuration

IN THIS SECTION

- Procedure | 290

Procedure

CLI Quick Configuration

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

```
set system security-profile ls-design-profile idp-policy Recommended
```

Step-by-Step Procedure

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

To add a predefined IDP policy to a security profile for a user logical system:

1. Log in to the primary logical system as the primary administrator and enter configuration mode.

```
[edit]
admin@host> configure
admin@host#
```

2. Add the IDP policy to the security profile.

```
[edit system security-profile]
admin@host# set ls-design-profile idp-policy Recommended
```

Results

From configuration mode, confirm your configuration by entering the `show security idp` and `show system security-profile ls-design-profile` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
admin@host# show security idp
    idp-policy Recommended {
        ...
```

```
    }
[edit]
admin@host# show system security-profile ls-design-profile
    policy {
        ...
    }
    idp-policy Recommended;
logical-system ls-product-design;
```

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

Verification

IN THIS SECTION

- Verifying the Configuration | 291

Verifying the Configuration

Purpose

Verify the IDP policy assigned to the logical system.

Action

From operational mode, enter the `show security idp logical-system policy-association` command. Ensure that the IDP policy in the security profile that is bound to the logical system is correct.

```
admin@host> show security idp logical-system policy-association
Logical system         IDP policy
ls-product-design      Recommended
```

SEE ALSO

Example: Enabling IDP in a User Logical System Security Policy | 292

IDP in Logical Systems Overview | 274

Example: Enabling IDP in a User Logical System Security Policy

IN THIS SECTION

- Requirements | 292
- Overview | 292
- Configuration | 293
- Verification | 295

This example shows how to enable IDP in a security policy in a user logical system.

Requirements

Before you begin:

- Log in to the user logical system as the logical system administrator. See "User Logical Systems Configuration Overview" on page 51.
- From configuration mode, use the `show system security-profile <profile-name> idp-policy` command to see the security policy resources allocated to the logical system.
- Configure an IDP security policy for the user logical system as the primary administrator. See "Example: Configuring and Assigning a Predefined IDP Policy for a User Logical System" on page 288.

Overview

In this example, you configure the ls-product-design user logical system as shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System" on page 57.

You enable IDP in a security policy that matches any traffic from the ls-product-design-untrust zone to the ls-product-design-trust zone. Enabling IDP in a security policy directs matching traffic to be checked against the IDP rulebases.

NOTE: This example uses the IDP policy configured and assigned to the ls-product-design user logical system by the primary administrator in "Example: Configuring and Assigning a Predefined IDP Policy for a User Logical System" on page 288.

Configuration

IN THIS SECTION

- Procedure | 293

Procedure

CLI Quick Configuration

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

```
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust policy
enable-idp match source-address any
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust policy
enable-idp match destination-address any
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust policy
enable-idp match application any
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust policy
enable-idp then permit application-services idp
```

Step-by-Step Procedure

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

To configure a security policy to enable IDP in a user logical system:

1. Log in to the logical system as the user logical system administrator and enter configuration mode.

```
[edit]
lsdesignadmin1@host:ls-product-design>configure
lsdesignadmin1@host:ls-product-design#
```

2. Configure a security policy that matches traffic from the ls-product-design-untrust zone to the ls-product-design-trust zone.

```
[edit security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust]
lsdesignadmin1@host:ls-product-design# set policy enable-idp match source-address any
lsdesignadmin1@host:ls-product-design# set policy enable-idp match destination-address any
lsdesignadmin1@host:ls-product-design# set policy enable-idp match application any
```

3. Configure the security policy to enable IDP for matching traffic.

```
[edit security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust]
lsdesignadmin1@host:ls-product-design# set policy enable-idp then permit application-services
idp
```

Results

From configuration mode, confirm your configuration by entering the `show security policies` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this `show` command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
lsdesignadmin1@host:ls-product-design# show security policies
    from-zone ls-product-design-untrust to-zone ls-product-design-trust {
        policy enable-idp {
            match {
                source-address any;
                destination-address any;
                application any;
            }
            then {
```

```
                permit {
                    application-services {
                        idp;
                    }
                }
            }
        }
        ...
    }
```

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

Verification

IN THIS SECTION

- Verifying Attack Matches | 295

Verifying Attack Matches

Purpose

Verify that attacks are being matched in network traffic.

Action

From operational mode, enter the `show security idp attack table` command.

```
admin@host> show security idp attack table
IDP attack statistics:
  Attack name                                #Hits
  FTP:USER:ROOT                              1
```

SEE ALSO

Example: Configuring and Assigning a Predefined IDP Policy for a User Logical System | 288

IDP in Logical Systems Overview | 274

User Logical Systems Configuration Overview | 51

Example: Configuring an IDP Policy for a User Logical System

IN THIS SECTION

- Requirements | 296
- Overview | 296
- Configuration | 297
- Verification | 303

This example shows how to configure and assign an IDP policy to a user logical system. After assigning the IDP policy, the traffic is sent from client to check for the attack detection on the configured custom attack.

Requirements

This example uses the following hardware and software components:

- Junos OS Release 18.3R1 and later
- an SRX4200 device

Before you configure IDP policy on user logical system:

- Configure security zones. See "Example: Configuring Security Zones for a User Logical Systems" on page 167.

Overview

In this example, you configure a custom attack that is used in an IDP policy. The IDP policy is specified and enabled using a security policy configured in the user logical system.

Configuration

IN THIS SECTION

- Configuring a user logical system | 297
- Configuring a Custom Attack | 298
- Configuring an IDP Policy for the User Logical System | 300
- Enabling IDP in a Security Policy | 302

To configure IDP in a user logical system:

Configuring a user logical system

CLI Quick Configuration

Step-by-Step Procedure

To configure a user logical system:

1. Configure a user logical system.

```
[edit]
user@host# set logical-system LSYS1
```

2. Exit from the configuration mode and enter to the operational mode.

```
user@host# exit
```

3. Login as LSYS1 user to the user logical sytem and enter to configuration mode.

```
user@host> set cli logical-system LSYS1
user@host:LSYS1> edit
user@host:LSYS1#
```

Results

From configuration mode, confirm your configuration by entering the `show logical-systems` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show logical-systems
    LSYS1 {
    }
```

Configuring a Custom Attack

CLI Quick Configuration

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

```
set security idp custom-attack my-http severity info
set security idp custom-attack my-http attack-type signature protocol-binding application HTTP
set security idp custom-attack my-http attack-type signature context http-get-url
set security idp custom-attack my-http attack-type signature pattern .*test.*
set security idp custom-attack my-http attack-type signature direction any
```

Step-by-Step Procedure

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

To configure a custom attack object:

1. Log in to the user logical system as LSYS1 and enter configuration mode.

```
[edit]
user@host:LSYS1#
```

2. Create the custom attack object and set the severity level.

```
[edit security idp]
user@host:LSYS1# set custom-attack my-http severity info
```

3. Configure stateful signature parameters.

```
[edit security idp]
user@host:LSYS1# set custom-attack my-http attack-type signature protocol-binding application
HTTP
user@host:LSYS1# set custom-attack my-http attack-type signature context http-get-url
user@host:LSYS1# set custom-attack my-http attack-type signature pattern .*test.*
user@host:LSYS1# set custom-attack my-http attack-type signature direction any
```

Results

From configuration mode, confirm your configuration by entering the `show security idp custom-attack my-http` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host:LSYS1# show security idp custom-attack my-http
severity info;
    attack-type {
    signature {
        protocol-binding {
            application HTTP;
        }
        context http-get-url;
        pattern .*test.*;
        direction any;
    }
}
```

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

Configuring an IDP Policy for the User Logical System

CLI Quick Configuration

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

```
set security idp idp-policy idpengine rulebase-ips rule 1 match from-zone any
set security idp idp-policy idpengine rulebase-ips rule 1 match source-address any
set security idp idp-policy idpengine rulebase-ips rule 1 match to-zone any
set security idp idp-policy idpengine rulebase-ips rule 1 match destination-address any
set security idp idp-policy idpengine rulebase-ips rule 1 match application default
set security idp idp-policy idpengine rulebase-ips rule 1 match attacks custom-attacks my-http
set security idp idp-policy idpengine rulebase-ips rule 1 then action no-action
set security idp idp-policy idpengine rulebase-ips rule 1 then notification log-attacks
```

Step-by-Step Procedure

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

To configure an IDP policy:

1. Create the IDP policy and configure match conditions.

```
[edit security idp]
user@host:LSYS1# set idp-policy idpengine rulebase-ips rule 1 match from-zone any
user@host:LSYS1# set idp-policy idpengine rulebase-ips rule 1 match source-address any
user@host:LSYS1# set idp-policy idpengine rulebase-ips rule 1 match to-zone any
user@host:LSYS1# set idp-policy idpengine rulebase-ips rule 1 match destination-address any
user@host:LSYS1# set idp-policy idpengine rulebase-ips rule 1 match application default
user@host:LSYS1# set idp-policy idpengine rulebase-ips rule 1 match attacks custom-attacks my-
http
```

2. Configure actions for the IDP policy.

```
[edit security idp]
user@host:LSYS1# set idp-policy idpengine rulebase-ips rule 1 then action no-action
user@host:LSYS1# set idp-policy idpengine rulebase-ips rule 1 then notification log-attacks
```

Results

From configuration mode, confirm your configuration by entering the `show security idp idp-policy idpengine` and `show system security-profile master-profile` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host:LSYS1# show security idp idp-policy idpengine
rulebase-ips {
    rule 1 {
        match {
            from-zone any;
            source-address any;
            to-zone any;
            destination-address any;
            application default;
            attacks {
                custom-attacks my-http;
            }
        }
        then {
            action {
                no-action;
            }
            notification {
                log-attacks;
            }
        }
    }
}
```

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

Enabling IDP in a Security Policy

CLI Quick Configuration

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

```
set security policies from-zone z1 to-zone z2 policy p1 match source-address any
set security policies from-zone z1 to-zone z2 policy p1 match destination-address any
set security policies from-zone z1 to-zone z2 policy p1 match application any
set security policies from-zone z1 to-zone z2 policy p1 then permit application-services idp-
policy idpengine
```

Step-by-Step Procedure

To enable IDP in a security policy:

1. Create the security policy and configure match conditions.

```
[edit security policies from-zone z1 to-zone z2]
user@host:LSYS1# set policy p1 match source-address any
user@host:LSYS1# set policy p1 match destination-address any
user@host:LSYS1# set policy p1 match application any
```

2. Enable IDP.

```
[edit security policies from-zone z1 to-zone z2]
user@host:LSYS1# set policy p1 then permit application-services idp-policy idpengine
```

Results

From configuration mode, confirm your configuration by entering the `show security policies` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host:LSYS1# show security policies
```

```
from-zone z1 to-zone z2 {
    policy p1{
        match {
            source-address any;
            destination-address any;
            application any;
        }
        then {
            permit {
                application-services {
                    idp-policy idpengine;
                }
            }
        }
    }
}
```

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

Verification

IN THIS SECTION

- Verifying Attack Detection | 303

To send traffic and check for attack detection from user logical system:

Verifying Attack Detection

Purpose

Verify that attack detection is happening for the custom attack.

Action

From operational mode, enter the `show security idp attack table` command.

```
user@host:LSYS1> show security idp policies
 PIC : FPC 0 PIC 0:
```

```
ID    Name                  Sessions    Memory      Detector
 1     idpengine             0           188584      12.6.130180509
```

```
user@host:LSYS1> show security idp attack table
IDP attack statistics:

  Attack name                                  #Hits
  my-http                                          1
```

Meaning

The output displays the attacks detected for the custom attack that is configured in the IDP policy in the user logical system LSYS1.

SEE ALSO

idp (logical-systems) | **796**

ALG for Logical Systems

IN THIS SECTION

- Understanding Application Layer Gateway (ALG) in Logical Systems | **305**
- Enabling and Disabling ALG for Logical System | **306**
- Example: Enabling FTP ALG in a Logical System | **311**

An Application Layer Gateway (ALG) in logical systems enables the gateway to parse application layer payloads and take decisions whether to allow or deny traffic to the application server. ALGs supports the applications such as Transfer Protocol (FTP) and various IP protocols that use the application layer payload to communicate the dynamic Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) ports on which the applications open data connections. For more information, see the following topics:

Understanding Application Layer Gateway (ALG) in Logical Systems

The primary administrator can configure ALGs at the root level. The configuration is inherited by all user logical systems. ALGs can also be configured discretely for user logical systems. The ALG status is not inherited by all user logical systems. For a newly created logical system, the ALG consists of a default status. The FTP protocol ALG can be enabled or disabled for a specific logical system. The ICMP ALG protocol is enabled by default and is not provisioned to disable.

NOTE: When an SRX device is upgraded to 18.2 release, the ALG status in a logical system is changed when compared with previous status. This change affects the ALG traffic in the logical system. For example, before upgrade, H.323 ALG is configured to enable by root. So H.323 ALG is also enabled in lsys1. After upgrade to 18.2, H.323 ALG status in lsys1 is disabled because the default status for H.323 is disabled for a new logical system.

NOTE: You can enable a particular ALG for only one specific logical system.

By default, the following ALGs are enabled on a root logical system:

- DNS
- FTP
- MSRPC
- PPTP
- SUNRPC
- TALK
- TFTP

Starting in Junos OS Release 18.2R1, you can either enable or disable the ALGs configuration for each logical system individually, and view the status of the ALGs for all logical systems or specific logical system. All 12 data ALGs (DNS, FTP, TFTP, MSRPC, SUNRPC, PPTP, RSH, RTSP, TALK, SQL, IKE, and TWAMP) and four VOIP ALGs (SIP, H.323, MGCP, and SCCP) are supported on the logical systems.

SEE ALSO

show security alg status logical-system | 989

Example: Enabling FTP ALG in a Logical System | 311

alg

Enabling and Disabling ALG for Logical System

This topic shows how to enable or disable the ALG status for each logical system.

1. By Default IKE ALG is disabled on the logical system. To enable this ALG, use the following command.
   - Enable IKE and ESP ALG with NAT.

```
[edit]
user@host# set logical-systems LSYS1 security alg ike-esp-nat enable
```

2. By default, the DNS, FTP, PPTP, SIP, SUNRPC and TWAMP ALGs are enabled on the logical system. To disable these ALGs, use the following commands.
   - Disable DNS ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg dns disable
```

   - Disable FTP ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg ftp disable
```

   - Disable H323 ALG.

```
[edit]
user@host# logical-systems LSYS1 security alg h323 disable
```

   - Disable MGCP ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg mgcp disable
```

- Disable MSRPC ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg msrpc disable
```

- Disable PPTP ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg pptp disable
```

- Disable RSH ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg rsh disable
```

- Disable RTSP ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg rtsp disable
```

- Disable SCCP ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg sccp disable
```

- Disable SIP ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg sip disable
```

- Disable SQL ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg sql disable
```

- Disable SUNRPC ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg sunrpc disable
```

- Disable TALK ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg talk disable
```

- Disable TFTP ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg tftp disable
```

3. Configuring ALG functions in logical systems.

- Configure DNS ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg dns
```

- Configure FTP ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg ftp
```

- Configure H323 ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg h323
```

- Configure IKE and ESP ALG with NAT.

```
[edit]
user@host# set logical-systems LSYS1 security alg ike-esp-nat
```

- Configure MGCP ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg mgcp
```

- Configure MSRPC ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg msrpc
```

- Configure PPTP ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg pptp
```

- Configure RSH ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg rsh
```

- Configure RTSP ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg rtsp
```

- Configure SCCP ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg sccp
```

- Configure SIP ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg sip
```

- Configure SQL ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg sql
```

- Configure SUNRPC ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg sunrpc
```

- Configure TALK ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg talk
```

- Configure TFTP ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg tftp
```

- Configure TWAMP ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg twamp
```

- Configure extended function for FTP ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg ftp allow-mismatch-ip-address
```

- Configure extended function for MSRPC ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg msrpc map-entry-timeout 10
```

- Configure extended function for SUNRPC ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg sunrpc map-entry-timeout 10
```

- Configure extended function for SIP ALG.

```
[edit]
user@host# set logical-systems LSYS1 security alg sip retain-hold-resource
```

Example: Enabling FTP ALG in a Logical System

IN THIS SECTION

- Requirements | 311
- Overview | 311
- Configuration | 312
- Verification | 318

This example shows how to enable or disable an FTP ALG configuration in a logical system and send traffic based on FTP ALG configuration of the logical system individually.

Requirements

Before you begin:

- Log in to the user logical system as the logical system administrator. See "User Logical Systems Configuration Overview" on page 51.

Overview

In this example, the ALG for FTP is configured to monitor and allow FTP traffic to be exchanged between the clients and the server on a logical system.

By default, the FTP ALG is enabled on the logical system.

Configuration

IN THIS SECTION

- CLI Quick Configuration | 312
- Configuring FTP ALG in a Logical System | 313
- Results | 316

CLI Quick Configuration

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

```
set system security-profile p1 policy maximum 100
set system security-profile p1 policy reserved 50
set system security-profile p1 zone maximum 100
set system security-profile p1 zone reserved 50
set system security-profile p1 flow-session maximum 6291456
set system security-profile p1 flow-session reserved 50
set system security-profile p1 flow-gate maximum 524288
set system security-profile p1 flow-gate reserved 50
set logical-systems LSYS0 interfaces lt-0/0/0 unit 0 encapsulation ethernet-vpls
set logical-systems LSYS0 interfaces lt-0/0/0 unit 0 peer-unit 1
set logical-systems LSYS0 routing-instances vr0 instance-type vpls
set logical-systems LSYS0 routing-instances vr0 interface lt-0/0/0.0
set system security-profile p1 logical-system LSYS0
set system security-profile p1 logical-system LSYS1
set logical-systems LSYS1 interfaces lt-0/0/0 unit 1 encapsulation ethernet
set logical-systems LSYS1 interfaces lt-0/0/0 unit 1 peer-unit 0
set logical-systems LSYS1 interfaces lt-0/0/0 unit 1 family inet address 10.0.0.0/8
set logical-systems LSYS1 interfaces ge-0/0/0 unit 0 family inet address 198.51.100.0/24
set logical-systems LSYS1 interfaces ge-0/0/1 unit 0 family inet address 203.0.113.0/24
set logical-systems LSYS1 security zones security-zone LSYS1_tzone host-inbound-traffic system-
services all
set logical-systems LSYS1 security zones security-zone LSYS1_tzone host-inbound-traffic protocol
all
set logical-systems LSYS1 security zones security-zone LSYS1_tzone interfaces ge-0/0/0
```

```
set logical-systems LSYS1 security zones security-zone LSYS1_utzone host-inbound-traffic system-
services all
set logical-systems LSYS1 security zones security-zone LSYS1_utzone host-inbound-traffic
protocol all
set logical-systems LSYS1 security zones security-zone LSYS1_utzone interfaces ge-0/0/1
set logical-systems LSYS1 security policies from-zone LSYS1_tzone to-zone LSYS1_utzone policy
p11 match source-address any
set logical-systems LSYS1 security policies from-zone LSYS1_tzone to-zone LSYS1_utzone policy
p11 match destination-address any
set logical-systems LSYS1 security policies from-zone LSYS1_tzone to-zone LSYS1_utzone policy
p11 match application junos-ftp
set logical-systems LSYS1 security policies from-zone LSYS1_tzone to-zone LSYS1_utzone policy
p11 match application junos-ping
set logical-systems LSYS1 security policies from-zone LSYS1_tzone to-zone LSYS1_utzone policy
p11 then permit
set logical-systems LSYS1 security policies default-policy deny-all
```

Configuring FTP ALG in a Logical System

Step-by-Step Procedure

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

To configure an ALG in a user logical system:

1. Configure a security profile.

```
[edit system security-profile]
user@host#set p1 policy maximum 100
user@host#set p1 policy reserved 50
user@host#set p1 zone maximum 100
user@host#set p1 zone reserved 50
user@host#set p1 flow-session maximum 6291456
user@host#set p1 flow-session reserved 50
user@host#set p1 flow-gate maximum 524288
user@host#set p1 flow-gate reserved 50
```

2. Configure the primary logical system.

Step-by-Step Procedure

a. Create the primary logical system

```
[edit logical-systems]
user@host#set LSYS0
user@host#set LSYS1
```

b. Configure interfaces for a primary logical system and configure logical tunnel interfaces and routing instances to the LSYS0.

```
[edit interfaces]
user@host#set lt-0/0/0 unit 0 encapsulation ethernet-vpls
user@host#set lt-0/0/0 unit 0 peer-unit 1
user@host#set routing-instances vr0 instance-type vpls
user@host#set routing-instances vr0 interface lt-0/0/0.0
```

c. Configure a security profile p1 and assign it to the root logical system LSYS0.

```
[edit system security-profile]
user@host#set p1 logical-system LSYS0
```

3. Configure a user logical system.

Step-by-Step Procedure

a. Create the user logical system LSYS1

```
[edit logical-systems]
user@host#set LSYS1
```

b. Configure user logical and logical tunnel interfaces to transfer traffic within the logical system.

```
[edit interfaces]
user@host#set ge-0/0/0 unit 0 family inet address 198.51.100.0/24
user@host#set ge-0/0/1 unit 0 family inet address 203.0.113.0/24
user@host#set lt-0/0/0 unit 1 encapsulation ethernet
```

```
user@host#set lt-0/0/0 unit 1 peer-unit 0
user@host#set lt-0/0/0 unit 1 family inet address 10.0.0.0/8
```

c. Assign a security profile p1 to LSYS1.

```
[edit system security-profile]
user@host#set p1 logical-system LSYS1
```

d. Configure security zones and assign interfaces to each zone.

```
[edit security zones]
user@host#set security-zone LSYS1_tzone host-inbound-traffic system-services all
user@host#set security-zone LSYS1_tzone host-inbound-traffic protocol all
user@host#set security-zone LSYS1_tzone interfaces ge-0/0/0
user@host#set security-zone LSYS1_utzone host-inbound-traffic system-services all
user@host#set security-zone LSYS1_utzone host-inbound-traffic protocol all
user@host#set security-zone LSYS1_utzone interfaces ge-0/0/1
```

4. Configure a security policy that permits FTP traffic from the LSYS1_tzone to LSYS1_utzone.

```
[edit security policies]
user@host#set from-zone LSYS1_tzone to-zone LSYS1_utzone policy p11 match source-address any
user@host#set from-zone LSYS1_tzone to-zone LSYS1_utzone policy p11 match destination-address
any
user@host#set from-zone LSYS1_tzone to-zone LSYS1_utzone policy p11 match application junos-
ftp
user@host#set from-zone LSYS1_tzone to-zone LSYS1_utzone policy p11 match application junos-
ping
user@host#set from-zone LSYS1_tzone to-zone LSYS1_utzone policy p11 then permit
user@host#set default-policy deny-all
```

Results

From configuration mode, confirm the configuration for LSYS0 and LSYS1 by entering the `show logical-systems`. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host#show logical-systems LSYS0
interfaces {
    lt-0/0/0 {
        unit 0 {
            encapsulation ethernet-vpls;
            peer-unit 1;
        }
        unit 2 {
            encapsulation ethernet-vpls;
            peer-unit 3;
        }
    }
}
routing-instances {
    vr0 {
        instance-type vpls;
        interface lt-0/0/0.0;
        interface lt-0/0/0.2;
    }
}
```

```
user@host#show logical-systems LSYS1
interfaces {
    lt-0/0/0 {
        unit 1 {
            encapsulation ethernet;
            peer-unit 0;
            family inet {
                address 10.0.1.1/24;
            }
        }
    }
    reth0 {
        unit 0 {
            family inet {
```

```
                address 198.51.100.0/24;
            }
        }
    }
}
security {
    alg{
        ftp;
    }
    policies {
        from-zone LSYS1_tzone to-zone LSYS1_utzone {
            policy P11 {
                match {
                    source-address any;
                    destination-address any;
                    application [ junos-ping junos-ftp ];
                }
                then {
                    permit;
                }
            }
        }
        default-policy {
            deny-all;
        }
    }
    zones {
        security-zone LSYS1_tzone {
            host-inbound-traffic {
                system-services {
                    all;
                }
                protocols {
                    all;
                }
            }
            interfaces {
                reth0.0;
            }
        }
        security-zone LSYS1_utzone {
            host-inbound-traffic {
                system-services {
```

```
                    all;
                }
                protocols {
                    all;
                }
            }
            interfaces {
                lt-0/0/0.1;
            }
        }
    }
}
```

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

Verification

IN THIS SECTION

- Verify ALG status for user logical system | **318**
- Verify ALG status for all the logical systems | **319**
- Verifying Intra-Logical System Traffic on a Logical System | **322**

To confirm that the configuration is working properly, perform these tasks:

Verify ALG status for user logical system

Purpose

Verify alg status for FTP is enabled.

Action

To verify the configuration is working properly, enter the `show security alg status logical-system LSYS1` command.

```
user@host> show security alg status logical-system LSYS1
ALG Status:
```

```
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Enabled
  SQL      : Disabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled
  TWAMP    : Disabled
```

Meaning

The output displays the alg status for FTP Enabled for the logical system LSYS1.

Verify ALG status for all the logical systems

Purpose

Verify the ALG status for all the logical systems on the device.

Action

To verify the configuration is working properly, enter the `show security alg status logical-system all` command.

```
user@host> show security alg status logical-system all
Logical system: root-logical-system
ALG Status:
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
```

```
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Disabled
  SQL      : Disabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled
  TWAMP    : Disabled

Logical system: LSYS3
ALG Status:
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Enabled
  SQL      : Disabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled
  TWAMP    : Disabled

Logical system: LSYS1
ALG Status:
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Enabled
  SQL      : Disabled
```

```
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled
  TWAMP    : Disabled

Logical system: LSYS2
ALG Status:
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Enabled
  SQL      : Disabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled
  TWAMP    : Disabled

Logical system: LSYS0
ALG Status:
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Disabled
  SQL      : Disabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled
```

```
   TWAMP    : Disabled
```

Meaning

The output displays the ALG status for all the logical systems on the device.

Verifying Intra-Logical System Traffic on a Logical System

Purpose

Verify the information about active resources, clients, groups, and sessions created through the resource manager.

Action

From operational mode, enter the `show security resource-manager summary` command.

```
 user@host> show security resource-manager summary
Active resource-manager clients   : 16
Active resource-manager groups    : 3
Active resource-manager resources : 26
Active resource-manager sessions  : 4
```

Meaning

The output displays summary information about active resources, clients, groups, and sessions created through the resource manager.

SEE ALSO

Understanding Application Layer Gateway (ALG) in Logical Systems | 305

show security alg status logical-system | 989

RELATED DOCUMENTATION

Security Zones in Logical Systems | 149

DHCP for Logical Systems

IN THIS SECTION

- Understanding DHCP Support for Logical Systems | 323
- Minimum DHCPv6 Relay Agent Configuration for Logical Systems | 323
- Example: Configuring the DHCPv6 Client for Logical Systems | 325
- Example: Configuring the DHCPv6 Server Options for Logical Systems | 333

Understanding DHCP Support for Logical Systems

Starting in Junos OS Release 18.4R1, a logical system supports the DHCP client feature to learn IP addresses for interfaces assigned to the logical systems. Additionally, starting in Junos OS Release 18.4R1, logical systems support the DHCP relay feature. A DHCP relay agent forwards DHCP requests and responses between the DHCP client and the DHCP server.

A DHCP server allocates IP addresses and provides IP configuration settings such as the DNS server and default gateway to client hosts on a subnet served by an interface of a logical system. The DHCP allows network administrators centrally manage a pool of IP addresses among hosts and automate the assignment of IP addresses in a network within a logical system. An IP address is leased to a host for a limited time period, allowing the DHCP server to share a limited IP addresses among a group of hosts that do not require permanent IP addresses.

An interface of an SRX Series device operating as a DHCP client receives the TCP or IP settings and the IP address from an external DHCP server.

An SRX Series device operating as a DHCP relay agent for logical systems forwards incoming requests from the DHCP clients to a specified DHCP server. The client requests pass through interfaces on the logical systems.

Minimum DHCPv6 Relay Agent Configuration for Logical Systems

The following example describes the minimum configuration required to configure an SRX Series device as a DHCPv6 relay agent for the logical system.

Before you begin determine the following:

- The DHCPv6 relay group and the DHCP active server-group for logical system.

1. Configure an interface with an IPv6 address for the logical system.

```
user@host# set logical-systems LSYS1 interfaces ge-0/0/0 unit 0 family inet6 address
2001:db8::1/64
```

2. Specify the name of the server-group and add the IP address for the DHCP servers belonging to the same group.

```
user@host# set logical-systems LSYS1 forwarding-options dhcp-relay dhcpv6 group inf interface
ge-0/0/0.0
```

3. Specify the name of the active server group.

```
user@host# set logical-systems LSYS1 forwarding-options dhcp-relay dhcpv6 active-server-group
server6
```

4. Create a DHCP relay group that includes at least one interface for the logical system.

```
user@host# set logical-systems LSYS1 forwarding-options dhcp-relay dhcpv6 server-group
server6 2001:db8::1/64
```

5. Confirm your configuration by entering the `show logical-systems LSYS1` command.

```
[edit]
user@host# show logical-systems LSYS1
interfaces {
    ge-0/0/0 {
        unit 0 {
            family inet6 {
                address 2001:db8::1/64;;
            }
        }
    }
}
forwarding-options {
    dhcp-relay {
```

```
        dhcpv6 {
            group inf {
                interface ge-0/0/0.0;
            }
            server-group {
                server6 {
                    2001:db8::1/64;
                }
            }
            active-server-group server6;
        }
    }
}
```

NOTE: To configure the DHCP relay agent in a routing instance for the logical system, configure the `dhcp-relay` statement in the `edit logical-systems LSYS1 routing-instances R1` hierarchy level.

Example: Configuring the DHCPv6 Client for Logical Systems

IN THIS SECTION

- Requirements | 325
- Overview | 326
- Configuration | 326
- Verification | 330

This example shows how to configure an SRX Series device as a DHCPv6 client for the logical systems.

Requirements

This example uses the following hardware and software components:

- An SRX Series device
- Junos OS Release 18.4R1

Before you begin:

- Read the Understanding DHCP Support for Logical Systems to understand how and where this procedure fits in the overall support for DHCP.

  No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, the primary administrator configures an SRX device as a DHCPv6 client for a logical system.

The DHCPv6 client for a logical system includes the following features:

- Identity association for non-temporary addresses (IA_NA)
- Identity association for prefix delegation (IA_PD)
- Autoconfig or stateful mode
- DHCP unique identifier (DUID)

Configuration

IN THIS SECTION

- CLI Quick Configuration | 326
- Configuring DHCPv6 Client in a Logical System | 327
- Procedure | 327
- Results | 329

CLI Quick Configuration

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

```
set logical-systems LSYS1 security zones security-zone trust host-inbound-traffic system-
services all
set logical-systems LSYS1 security zones security-zone trust host-inbound-traffic protocols all
set logical-systems LSYS1 security zones security-zone trust interfaces ge-0/0/0.0
```

```
set logical-systems LSYS1 routing-instances r1 instance-type virtual-router
set logical-systems LSYS1 routing-instances r1 interface ge-0/0/0.0
set logical-systems LSYS1 interfaces ge-0/0/0 unit 0 family inet6 dhcpv6-client client-type
autoconfig
set logical-systems LSYS1 interfaces ge-0/0/0 unit 0 family inet6 dhcpv6-client client-type
stateful
set logical-systems LSYS1 interfaces ge-0/0/0 unit 0 family inet6 dhcpv6-client client-ia-type
ia-na
set logical-systems LSYS1 interfaces ge-0/0/0 unit 0 family inet6 dhcpv6-client client-ia-type
ia-pd
set logical-systems LSYS1 interfaces ge-0/0/0 unit 0 family inet6 dhcpv6-client client-
identifier duid-type duid-ll
set logical-systems LSYS1 interfaces ge-0/0/0 unit 0 family inet6 dhcpv6-client req-option dns-
server
set protocols router-advertisement interface ge-0/0/0.0
```

Configuring DHCPv6 Client in a Logical System

Procedure

Step-by-Step Procedure

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

1. Configure the security zones to permit traffic for a logical system.

```
[edit logical-systems LSYS1 security zones]
user@host# set security-zone trust host-inbound-traffic system-services all
user@host# set security-zone trust host-inbound-traffic protocols all
user@host# set security-zone trust interfaces ge-0/0/0.0
```

2. Create a routing instance and assign the routing instance type for a logical system.

```
[edit logical-systems LSYS1]
user@host# set routing-instances r1 instance-type virtual-router
```

3. Specify the interface name for the routing instance.

```
[edit logical-systems LSYS1]
user@host# set routing-instances r1 interface ge-0/0/0.0
```

4. Configure the DHCPv6 client type. The client type can be `autoconfig` or `stateful` for the logical system.

   - To enable the DHCPv6 auto configuration mode, configure the client type as `autoconfig`.

```
[edit logical-systems LSYS1 interfaces ge-0/0/0 unit 0 family inet6 dhcpv6-client]
user@host# set client-type autoconfig
```

   - For stateful address assignment, configure the client type as `stateful`.

```
[edit logical-systems LSYS1 interfaces ge-0/0/0 unit 0 family inet6 dhcpv6-client]
user@host# set client-type stateful
```

5. Specify the identity association type.

   - To configure identity association for nontemporary address (IA_NA) assignment, specify the `client-ia type` as `ia-na`.

```
[edit logical-systems LSYS1 interfaces ge-0/0/0 unit 0 family inet6 dhcpv6-client]
user@host# set client-ia-type ia-na
```

   - To configure identity association for prefix delegation (IA_PD), specify the `client-ia-type` as `ia-pd`.

```
[edit logical-systems LSYS1 interfaces ge-0/0/0 unit 0 family inet6 dhcpv6-client]
user@host# set client-ia-type ia-pd
```

6. Configure the DHCPv6 client identifier by specifying the DHCP unique identifier (DUID) type for the logical system. The following DUID type is supported:

   - Link layer address (duid-ll)

```
[edit tlogical-systems LSYS1 interfaces ge-0/0/0 unit 0 family inet6 dhcpv6-client]
user@host# set client-identifier duid-type duid-ll
```

7. Specify the DHCPv6 client requested option as `dns-server` for the logical system.

```
[edit logical-systems LSYS1 interfaces ge-0/0/0 unit 0 family inet6 dhcpv6-client]
user@host# set req-option dns-server
```

8. Configure the router advertisement.

```
[edit]
user@host# set protocols router-advertisement interface ge-0/0/0.0
```

Results

- From configuration mode, confirm your configuration by entering the `show logical-systems LSYS1` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show logical-systems LSYS1
interfaces {
    ge-0/0/0 {
        unit 0 {
            family inet6 {
                dhcpv6-client {
                    client-type stateful;
                    client-ia-type ia-na;
                    client-ia-type ia-pd;
                    client-identifier duid-type duid-ll;
                    req-option dns-server;
                }
            }
        }
    }
}
routing-instances {
    r1 {
        instance-type virtual-router;
        interface ge-0/0/0.0;
    }
}
security {
    zones {
```

```
        security-zone trust {
            host-inbound-traffic {
                system-services {
                    all;
                }
                protocols {
                    all;
                }
            }
            interfaces {
                ge-0/0/0.0;
            }
        }
    }
}
```

- From configuration mode, confirm your configuration by entering the `show protocols` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show protocols
router-advertisement {
    interface ge-0/0/0.0;
}
```

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

Verification

IN THIS SECTION

- Verifying the DHCPv6 Client for Logical Systems | **331**
- Verifying the DHCPv6 Client Binding for Logical Systems | **331**
- Verifying the DHCPv6 Client Statistics for Logical Systems | **332**

To confirm that the configuration is working properly, perform these tasks:

Verifying the DHCPv6 Client for Logical Systems

Purpose

Verify that the DHCPv6 client information is configured.

Action

From the operational mode, enter the `show dhcpv6 client binding logical-systems LSYS1` command.

```
user@host> show dhcpv6 client binding  logical-systems LSYS1
 IP/prefix                       Expires     State      ClientType    Interface       Client DUID
        2000::17/128                    67762       BOUND      STATEFUL      ge-0/0/6.0
LL0x3-10:0e:7e:49:25:86
        2000:100::/64                   67762       BOUND      STATEFUL      ge-0/0/6.0
LL0x3-10:0e:7e:49:25:86
```

Meaning

The output displays the address binding information for the logical system.

Verifying the DHCPv6 Client Binding for Logical Systems

Purpose

Verify that the DHCPv6 client binding information is configured.

Action

From the operational mode, enter the `show dhcpv6 client binding detail logical-systems LSYS1` command.

```
user@host> show dhcpv6 client binding detail logical-systems LSYS1
Client Interface/Id: ge-0/0/6.0
     Hardware Address:             10:0e:7e:49:25:86
     State:                        BOUND(DHCPV6_CLIENT_STATE_BOUND)
     ClientType:                   STATEFUL
     Lease Expires:                2018-11-09 07:11:47 UTC
     Lease Expires in:             67760 seconds
     Lease Start:                  2018-11-08 07:11:47 UTC
```

```
     Bind Type:                    IA_NA IA_PD
     Preferred prefix length       0
     Sub prefix length             0
     Client DUID:                  LL0x3-10:0e:7e:49:25:86
     Rapid Commit:                 Off
     Server Identifier:            fe80::46f4:77ff:fed6:670a
     Client IP Address:            2000::17/128
     Client IP Prefix:             2000:100::/64


     DHCP options:
    Name: server-identifier, Value: VENDOR0x00000583-0x34343a34
```

Meaning

The output displays the detailed client binding information for the logical system.

Verifying the DHCPv6 Client Statistics for Logical Systems

Purpose

Verify that the DHCPv6 client statistics information is configured.

Action

From the operational mode, enter the `show dhcpv6 client statistics logical-systems LSYS1` command.

```
user@host> show dhcpv6 client statistics logical-systems LSYS1
Dhcpv6 Packets dropped:
     Total               3
     Bad Send            3

    Messages received:
    DHCPV6_ADVERTISE           1
    DHCPV6_REPLY               1
    DHCPV6_RECONFIGURE         0

    Messages sent:
    DHCPV6_DECLINE             0
    DHCPV6_SOLICIT             1
```

```
    DHCPV6_INFORMATION_REQUEST 0
    DHCPV6_RELEASE             0
    DHCPV6_REQUEST             1
    DHCPV6_CONFIRM             0
    DHCPV6_RENEW               0
    DHCPV6_REBIND              0
```

Meaning

The output displays the information about the number of packets discarded, the number of messages received and the number of messages sent by the DHCP client for the logical system.

Example: Configuring the DHCPv6 Server Options for Logical Systems

IN THIS SECTION

- Requirements | 333
- Overview | 334
- Configuration | 334
- Verification | 338

This example shows how to configure DHCPv6 server options on SRX Series devices for the logical system.

Requirements

This example uses the following hardware and software components:

- An SRX Series device
- Junos OS Release 18.4R1

Before you begin determine the following:

- The IPv6 address pool range and the IPv6 prefix for logical systems.

Overview

In this example, you set a default client limit as 200 for all DHCPv6 groups. You then create a group called `my-group` that contains at least one interface. In this case, the interface is ge-0/0/2.0. You set a range of interfaces using the `upto` command and set a custom client limit as 200 for group `my-group` that overrides the default limit. Finally, you configure interface ge-0/0/2.0 with IPv6 address 2001:db8::1/64 and set router advertisement for interface ge-0/0/2.0.

Configuration

IN THIS SECTION

- CLI Quick Configuration | 334
- Procedure | 335
- Results | 336

CLI Quick Configuration

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

```
set logical-systems LSYS1 system services dhcp-local-server dhcpv6 group my-group overrides
interface-client-limit 200
set logical-systems LSYS1 system services dhcp-local-server dhcpv6 group my-group interface
ge-0/0/2.0
set logical-systems LSYS1 interfaces ge-0/0/2 unit 0 family inet6 address 2001:db8::1/64
set logical-systems LSYS1 access address-assignment pool my-pool family inet6 prefix
2001:db8::1/64
set logical-systems LSYS1 access address-assignment pool my-pool family inet6 range r1 low
2001:db8::1/64
set logical-systems LSYS1 access address-assignment pool my-pool family inet6 range r1 high
2001:db8::1/64
set logical-systems LSYS1 access address-assignment pool my-pool family inet6 dhcp-attributes
maximum-lease-time 200
set logical-systems LSYS1 access address-assignment pool my-pool family inet6 dhcp-attributes
option 21 string sip1.net
```

```
set logical-systems LSYS1 protocols router-advertisement interface ge-0/0/2.0 prefix
2001:db8::1/64
```

Procedure

Step-by-Step Procedure

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

To configure the DHCPv6 server options for logical systems:

1. Configure a DHCP local server.

```
[edit logical-systems LSYS1]
user@host# set system services dhcp-local-server dhcpv6
```

2. Set a default limit for all DHCPv6 groups.

```
[edit logical-systems LSYS1 system services dhcp-local-server dhcpv6]
user@host# set group my-group overrides interface-client-limit 200
```

3. Specify a group name and interface.

```
[edit logical-systems LSYS1 system services dhcp-local-server dhcpv6]
user@host# set group my-group interface ge-0/0/2.0
```

4. Configure an interface with an IPv6 address.

```
[edit logical-systems LSYS1 interfaces]
user@host# set ge-0/0/2 unit 0 family inet6 address 2001:db8::1/64
```

5. Configure an address-pool and specify the IPv6 family.

```
[edit logical-systems LSYS1 access]
user@host# set address-assignment pool my-pool family inet6 prefix 2001:db8::1/64
```

6. Configure the IPv6 prefix, the range name, and the IPv6 range for the DHCPv6 clients

```
[edit logical-systems LSYS1 access]
user@host# set address-assignment pool my-pool family inet6 range r1 low 2001:db8::1/64
user@host# set address-assignment pool my-pool family inet6 range r1 high 2001:db8::1/64
```

7. Configure the DHCPv6 attribute for the maximum lease time.

```
[edit logical-systems LSYS1 access]
user@host# set address-assignment pool my-pool family inet6 dhcp-attributes maximum-lease-
time 200
```

8. Configure the user-defined option.

```
[edit logical-systems LSYS1 access]
user@host# set address-assignment pool my-pool family inet6 dhcp-attributes option 21 string
sip1.net
```

9. Configure the router advertisement for the interface.

```
[edit logical-systems LSYS1 protocols]
user@host# set router-advertisement interface ge-0/0/2.0 prefix 2001:db8::1/64
```

Results

From configuration mode, confirm your configuration by entering the `show logical-systems LSYS1` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show logical-systems LSYS1
interfaces {
    ge-0/0/2 {
        unit 0 {
            family inet6 {
                address 2001:db8::1/64;
            }
        }
    }
```

```
}
protocols {
    router-advertisement {
        interface ge-0/0/2.0 {
            prefix 2001:db8::1/64;
        }
    }
}
system {
    services {
        dhcp-local-server {
            dhcpv6 {
                group my-group {
                    overrides {
                        interface-client-limit 200;
                    }
                    interface ge-0/0/2.0;
                }
            }
        }
    }
}
access {
    address-assignment {
        pool my-pool {
            family inet6 {
                prefix 2001:db8::1/64;
                range r1 {
                    low 2001:db8::1/64;
                    high 2001:db8::1/64;
                }
                dhcp-attributes {
                    maximum-lease-time 200;
                    option 21 string sip1.net;
                }
            }
        }
    }
}
```

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

Verification

IN THIS SECTION

- Verifying the DHCPv6 Local Server Configuration | 338

To confirm that the configuration is working properly, perform these tasks:

Verifying the DHCPv6 Local Server Configuration

Purpose

Displays the address bindings in the client table on the extended DHCPv6 local server.

Action

From operational mode, enter the `show dhcpv6 server binding summary` command to display the address bindings in the client table on the DHCPv6 local server.

```
user@host> show dhcpv6 server binding summary
5 clients, (0 init, 5 bound, 0 selecting, 0 requesting, 0 renewing, 0 releasing)
```

Meaning

The output displays the information about the DHCPv6 local server address binding summary.

Application Security in Logical Systems

IN THIS SECTION

- Understanding Logical Systems Application Identification Services | 339
- Understanding Logical Systems Application Firewall Services | 341

- Example: Configuring Application Firewall Services for a Primary Logical Systems | 342
- Understanding Logical Systems Application Tracking Services | 348
- Example: Configuring Application Firewall Services for a User Logical System | 349
- Example: Configuring AppTrack for a User Logical Systems | 355

Application Security in logical systems enables to identify application traffic traversing your network regardless of port, protocol, and encryption, thereby providing greater visibility to control network traffic. The application security controls network traffic by setting and enforcing security policies based on accurate application information. For more information, see the following topics:

Understanding Logical Systems Application Identification Services

Predefined and custom application signatures identify an application by matching patterns in the first few packets of a session. Identifying applications provides the following benefits:

- Allows Intrusion Detection and Prevention (IDP) to apply appropriate attack objects to applications running on nonstandard ports.
- Improves performance by narrowing the scope of attack signatures for applications without decoders.
- Enables you to create detailed reports using AppTrack on applications passing through the device.

With logical systems, predefined and custom application signatures are global resources that are shared by all logical systems. The primary administrator is responsible for downloading and installing predefined Juniper Networks application signatures and creating custom application and nested application signatures to identify applications that are not part of the predefined database.

Application identification is enabled by default.

The application system cache (ASC) saves the mapping between an application type and the corresponding destination IP address, destination port, protocol type, and service. Each user logical system has its own ASC. A user logical system administrator can display the ASC entries for their logical system with the `show services application-identification application-system-cache` command. A user logical system administrator can use the `clear services application-identification application-system-cache` command to clear the ASC entries for their logical system.

Starting in Junos OS Release 18.2R1, the default behavior of the ASC is changed as follows:

- Security services including security policies, application firewall (AppFW), application tracking (AppTrack), application quality of service (AppQoS), Juniper Sky ATP, IDP, and UTM do not use the ASC by default.
- Miscellaneous services including advanced policy-based routing (APBR) use the ASC for application identification by default.

  For more information, see *Enabling or Disabling Application System Cache for Application Services*.

The primary administrator can display or clear ASC entries for any logical system. The primary administrator can also display or clear global counters with the `show services application-identification counter` and `clear services application-identification counter` commands.

Application signature package is installed at the global-level, that is shared by all user logical systems. The primary logical system administrator can install or uninstall application signature package.

Starting in Junos OS Release 18.3R1, the application identification (AppID) support for logical systems include two new options to view and clear logical system statistics and logical system counters statistics.

The primary logical system administrator can display or clear the statistics for all logical systems whereas the administrator for the user logical system can display or clear the statistics for their own logical system.

The user logical system administrator can view the AppID signature package status and version. Custom signatures configured by the primary logical system administrator can be configured in the use logical system security policies.

You can view the status and version information about the AppID signature package status and version by using the commands `show services application-identification status` and `show services application-identification version`.

SEE ALSO

Understanding the Junos OS Application Identification Database

Example: Scheduling the Application Signature Package Updates

Example: Configuring Junos OS Application Identification Custom Application Signatures

Understanding IDP Application Identification

Understanding the Application System Cache

Verifying Application System Cache Statistics

Understanding Logical Systems Application Firewall Services

An application firewall enables administrators of logical systems to create security policies for traffic based on application identification defined by application signatures. The application firewall provides additional security protection against dynamic-application traffic that might not be adequately controlled by standard network firewall policies. The application firewall controls information transmission by allowing or blocking traffic originating from particular applications.

To configure an application firewall, you define a rule set that contains rules specifying the action to be taken on identified dynamic applications. The rule set is configured independently and assigned to a security policy. Each rule set contains at least two rules, a matched rule (consisting of match criteria and action) and a default rule.

- A matched rule defines the action to be taken on matching traffic. When traffic matches an application and other criteria specified in the rule, the traffic is allowed or blocked based on the action specified in the rule.
- A default rule is applied when traffic does not match any other rule in the rule set.

The primary administrator can download a predefined application signature database from the Juniper Networks Security Engineering website or can define application signatures using the Junos OS configuration CLI. For more information about application identification and application signatures, see Application Security User Guide for Security Devices.

Configuring an application firewall on a logical system is the same process as configuring an application firewall on a device that is not configured with logical systems. However, the application firewall applies only to the logical system for which it is configured. The primary administrator can configure, enable, and monitor application firewalls on the primary logical system and all user logical systems on a device. User logical system administrators can configure, enable, and monitor application firewalls only on the user logical systems for which they have access.

SEE ALSO

Example: Configuring Application Firewall Services for a Primary Logical Systems | 342

Example: Configuring Application Firewall Services for a User Logical System | 349

Example: Configuring Application Firewall Services for a Primary Logical Systems

IN THIS SECTION

- Requirements | 342
- Overview | 343
- Configuration | 343
- Verification | 346

This example describes how to configure application firewall services on the primary, or root, logical system by a primary administrator. Only the primary administrator can configure, manage, and view configuration of the primary logical system, in addition to all user logical systems.

After configuring application firewall rule sets and rules, the primary administrator adds the application firewall rule set information to the security policy on the primary logical system.

For information about configuring an application firewall within a security policy, see *Application Firewall Overview*.

Requirements

Before you begin:

- Verify that all interfaces, routing instances, and security zones have been configured on the primary logical system.

  See "Example: Configuring Security Features for the Primary Logical Systems" on page 176.

- Verify that application firewall resources (appfw-rule-set and appfw-rule) have been allocated in a security profile and bound to the primary logical system through the [`system security-profile`] command. For application firewall resources, a security profile configuration allows 0 to 10,000 rule sets and 0 to 10,000 rules.

  NOTE: The primary administrator allocates various global system resources through a security profile configuration which is then bound to the various logical systems on the device. The primary administrator owns this function and configures the security profile for all user logical systems as well as the primary logical system.

For more information, see "Understanding Logical Systems Security Profiles (Primary Administrators Only)" on page 71.

- Log in to the primary logical system as the primary administrator.

  For information about primary administrator role functions, see "Understanding the Primary Logical Systems and the Primary Administrator Role" on page 23.

Overview

IN THIS SECTION

- Topology | 343

In this example you create application firewall services on the primary logical system, called root-logical-system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System" on page 57.

This example creates the following application firewall configuration:

- Rule set, root-rs1, with rules r1 and r2. When r1 is matched, telnet traffic is allowed through the firewall. When r2 is matched, web traffic is allowed through the firewall.
- Rule set, root-rs2, with rule r1. When r1 is matched, example2 traffic is blocked by the firewall.

All rule sets require a default rule, which specifies whether to permit or deny traffic that is not specified in any rules of a rule set. The default-rule action (permit or deny) must be the opposite from the action that is specified for the other rule(s) in the rule set.

Topology

Configuration

IN THIS SECTION

- Procedure | 344

Procedure

CLI Quick Configuration

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

```
set logical-systems root-logical-system security application-firewall rule-sets root-rs1 rule r1
match dynamic-application junos:telnet
set logical-systems root-logical-system security application-firewall rule-sets root-rs1 rule r1
then permit
set logical-systems root-logical-system security application-firewall rule-sets root-rs1 rule r2
match dynamic-application-group junos:web
set logical-systems root-logical-system security application-firewall rule-sets root-rs1 rule r2
then permit
set logical-systems root-logical-system security application-firewall rule-sets root-rs1 default-
rule deny
set logical-systems root-logical-system security application-firewall rule-sets root-rs2 rule r1
match dynamic-application junos:facebook
set logical-systems root-logical-system security application-firewall rule-sets root-rs2 rule r1
then deny
set logical-systems root-logical-system security application-firewall rule-sets root-rs2 default-
rule permit
```

Step-by-Step Procedure

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

To configure application firewall for a primary logical system:

1. Log in to the primary logical system as the primary administrator. See "Example: Configuring Root Password for Logical Systems" on page 55 and enter configuration mode.

```
admin@host> configure
admin@host#
```

2. Configure an application firewall rule set for root-logical-system.

```
[edit ]
admin@host# set logical-systems security application-firewall rule-sets root-rs1
```

3. Configure a rule for this rule set and specify which dynamic applications and dynamic application groups the rule should match.

```
[edit]
admin@host# set logical-systems security application-firewall rule-sets root-rs1 rule r1
match dynamic-application telnet then permit
```

4. Configure the default rule for this rule set and specify the action to take when the identified dynamic application is not specified in any rules of the rule set.

```
[edit]
admin@host# set logical-systems security application-firewall rule-sets root-rs1 default-rule
deny
```

5. Repeat these steps to configure another rule set, root-rs2, if desired.

Results

From configuration mode, confirm your configuration by entering the `show security application-firewall rule-sets` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this `show` command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
admin@host# show security application-firewall rule-sets all
...
    application-firewall {
        rule-sets root-rs1  {
            rule r1 {
                match {
                    dynamic-application [junos:telnet];
                }
```

```
                then {
                     permit;
                }
            }
            default-rule {
                deny;
            }
        }
        rule-sets root-rs1  {
            rule r2 {
                match {
                    dynamic-application-group [junos:web];
                }
                then {
                     permit;
                }
            }
            rule-sets root-rs2  {
                rule r1 {
                    match {
                        dynamic-application [junos:FACEBOOK];
                    }
                    then {
                        deny;
                    }
                }
                default-rule {
                    permit;
                }
            }
```

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

Verification

IN THIS SECTION

- Verifying Application Firewall Configuration | 347

Confirm that the configuration is working properly.

Verifying Application Firewall Configuration

Purpose

View the application firewall configuration on the primary logical system.

Action

From operational mode, enter the `show security application-firewall rule-set logical-system root-logical-system rule-set all` command.

```
admin@host> show security application-firewall rule-set logical-system root-logical-system rule-set all

Rule-set: root-rs1
    Logical system: root-logical-system
    Rule: r1
        Dynamic Applications: junos:telnet
        Action:permit
        Number of sessions matched: 10
Default rule:deny
        Number of sessions matched: 100
Number of sessions with appid pending: 2

Rule-set: root-rs1
    Logical system: root-logical-system
    Rule: r2
        Dynamic Applications: junos:web
        Action:permit
        Number of sessions matched: 20
Default rule:deny
        Number of sessions matched: 200
Number of sessions with appid pending: 4

Rule-set: root-rs2
    Logical system: root-logical-system
    Rule: r1
        Dynamic Applications: junos:FACEBOOK
        Action:deny
        Number of sessions matched: 40
Default rule:permit
```

```
        Number of sessions matched: 400
Number of sessions with appid pending: 10
```

SEE ALSO

- SRX Series Logical Systems Primary Administrator Configuration Tasks Overview | **24**
- Understanding Logical Systems Security Profiles (Primary Administrators Only) | **71**
- Understanding Logical Systems Application Firewall Services | **341**
- Example: Configuring Security Features for the Primary Logical Systems | **176**

Understanding Logical Systems Application Tracking Services

AppTrack is an application tracking tool that provides statistics for analyzing bandwidth usage of your network. When enabled, AppTrack collects byte, packet, and duration statistics for application flows in the specified zone. By default, when each session closes, AppTrack generates a message that provides the byte and packet counts and duration of the session, and sends it to the host device. The Security Threat Response Manager (STRM) retrieves the data and provides flow-based application visibility.

AppTrack can be enabled and configured within any logical system. Configuring AppTrack in a logical system is the same as configuring AppTrack on a device that is not configured for logical systems. An AppTrack configuration only applies to the logical system in which it is configured. The name of the logical system is added to AppTrack logs. The primary administrator can configure AppTrack for any logical system while a user logical system administrator can only configure AppTrack for the logical system that they are logged in to.

NOTE: The system log configuration is global on the device and must be configured by the primary administrator. The user logical system administrator cannot configure system logging for a logical system.

Counters keep track of the number of log messages sent and logs that have failed. AppTrack counters are global to the device. The primary administrator as well as user logical system administrators can view AppTrack counters with the `show security application-tracking counters` command.

SEE ALSO

- Understanding AppTrack

Example: Configuring AppTrack

Example: Configuring AppTrack for a User Logical Systems | 355

Example: Configuring Application Firewall Services for a User Logical System

IN THIS SECTION

- Requirements | 349
- Overview | 350
- Configuration | 350
- Verification | 353

This example describes how to configure application firewall services on a user logical system by a user logical system administrator. User logical system administrators can manage and monitor their own system application firewall rule sets and rules and manage the dynamic applications allowed or blocked on their respective logical systems.

After configuring application firewall rule sets and rules, user logical system administrators add the application firewall rule set information to the security policy on their individual logical systems.

For information about configuring an application firewall within a security policy, see *Application Firewall Overview*.

Requirements

Before you begin:

- Verify that the security zones are configured for the user logical system.
- Verify that the primary administrator has allocated application firewall resources (appfw-rule-set and appfw-rule) in the security profile bound to the user logical system.

  For more information, see "Understanding Logical Systems Security Profiles (Primary Administrators Only)" on page 71.
- Log in to the logical system as the user logical system administrator.

For information about user logical system administrator role functions, see "Understanding User Logical Systems and the User Logical System Administrator Role" on page 53.

Overview

IN THIS SECTION

- Topology | 350

In this example you configure application firewall services on the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System" on page 57.

This example creates the following application firewall configuration:

- Rule set, ls-product-design-rs1, with rules r1 and r2. When r1 is matched, telnet traffic is allowed through the firewall. When r2 is matched, web traffic is allowed through the firewall.
- Rule set, ls-product-design-rs2, with rule r1. When r1 is matched, Facebook traffic is blocked by the firewall.

All rule sets require a default rule, which specifies whether to permit or deny traffic that is not specified in any rules of a rule set. The default-rule action (permit or deny) must be the opposite from the action that is specified for the other rule(s) in the rule set.

Topology

Configuration

IN THIS SECTION

- Procedure | 351

Procedure

CLI Quick Configuration

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

```
set security application-firewall rule-sets ls-product-design-rs1 rule r1 match dynamic-
application junos:telnet
set security application-firewall rule-sets ls-product-design-rs1 rule r1 then permit
set security application-firewall rule-sets ls-product-design-rs1 rule r2 match dynamic-
application-group junos:web
set security application-firewall rule-sets ls-product-design-rs1 rule r2 then permit
set security application-firewall rule-sets ls-product-design-rs1 default-rule deny
set security application-firewall rule-sets ls-product-design-rs2 rule r1 match dynamic-
application junos:facebook
set security application-firewall rule-sets ls-product-design-rs2 rule r1 then deny
set security application-firewall rule-sets ls-product-design-rs2 default-rule permit
```

Step-by-Step Procedure

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

To configure application firewall for a user logical system:

1. Log in to the user logical system as the user logical system administrator and enter configuration mode.

```
lsdesignadmin1@host:ls-product-design> configure
lsdesignadmin1@host:ls-product-design#
```

2. Configure an application firewall rule set for this logical system.

```
[edit]
lsdesignadmin1@host:ls-product-design# set security application-firewall rule-sets ls-product-
design-rs1
```

3. Configure a rule for this rule set and specify which dynamic applications and dynamic application groups the rule should match.

```
[edit]
lsdesignadmin1@host:ls-product-design# set security application-firewall rule-sets ls-product-design-rs1 rule r1 match dynamic-application telnet then permit
```

4. Configure the default rule for this rule set and specify the action to take when the identified dynamic application is not specified in any rules of the rule set.

```
[edit]
lsdesignadmin1@host:ls-product-design# set security application-firewall rule-sets ls-product-design-rs1 default-rule deny
```

5. Repeat these steps to configure another rule set, ls-product-design-rs2, if desired.

Results

From configuration mode, confirm your configuration by entering the `show security application-firewall rule-set all` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this `show` command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
lsdesignadmin1@host:ls-product-design# show security application-firewall rule-set all
...
    application-firewall {
        rule-sets ls-product-design-rs1  {
            rule r1 {
                match {
                    dynamic-application [junos:telnet];
                }
                then {
                     permit;
                }
            }
            default-rule {
                deny;
```

```
            }
        }
        rule-sets ls-product-design-rs1  {
            rule r2 {
                match {
                    dynamic-application-group [junos:web];
                }
                then {
                     permit;
                }
            }
            rule-sets ls-product-design-rs2  {
                rule r1 {
                    match {
                        dynamic-application [junos:FACEBOOK];
                    }
                    then {
                        deny;
                    }
                }
                default-rule {
                    permit;
                }
            }
```

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

Verification

IN THIS SECTION

- Verifying Application Firewall Configuration | 354

Confirm that the configuration is working properly.

Verifying Application Firewall Configuration

Purpose

View the application firewall configuration on the user logical system.

Action

From operational mode, enter the `show security application-firewall rule-set all` command.

```
lsdesignadmin1@host:ls-product-design> show security application-firewall rule-set all

Rule-set: ls-product-design-rs1
    Logical system: ls-product-design
    Rule: r1
        Dynamic Applications: junos:telnet
        Action:permit
        Number of sessions matched: 10
Default rule:deny
        Number of sessions matched: 100
Number of sessions with appid pending: 2

Rule-set: ls-product-design-rs1
    Logical system: ls-product-design
    Rule: r2
        Dynamic Applications: junos:web
        Action:permit
        Number of sessions matched: 20
Default rule:deny
        Number of sessions matched: 200
Number of sessions with appid pending: 4

Rule-set: ls-product-design-rs2
    Logical system: ls-product-design
    Rule: r1
        Dynamic Applications: junos:FACEBOOK
        Action:deny
        Number of sessions matched: 40
Default rule:permit
        Number of sessions matched: 400
Number of sessions with appid pending: 10
```

SEE ALSO

User Logical Systems Configuration Overview | 51

Understanding Logical Systems Application Firewall Services | 341

Example: Configuring AppTrack for a User Logical Systems

IN THIS SECTION

- Requirements | 355
- Overview | 355
- Configuration | 356
- Verification | 358

This example shows how to configure the AppTrack tracking tool so you can analyze the bandwidth usage of your network.

Requirements

Before you begin:

- Log in to the user logical system as the logical system administrator. See "User Logical Systems Configuration Overview" on page 51.
- (Primary administrator) Configure system logging in the primary logical system. See Network Management and Monitoring Guide.

Overview

IN THIS SECTION

- Topology | 356

This example shows how to enable application tracking for the security zone ls-product-design-trust in the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System" on page 57.

The first message is generated at session start and update messages are sent every 5 minutes after that or until the session ends. A final message is sent at session end.

Topology

Configuration

IN THIS SECTION

- Procedure | 356

Procedure

CLI Quick Configuration

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

```
set security zones security-zone ls-product-design-trust application-tracking
set security application-tracking first-update
```

Step-by-Step Procedure

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

To configure AppTrack for a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

```
lsdesignadmin1@host:ls-product-design> configure
lsdesignadmin1@host:ls-product-design#
```

2. Enable AppTrack for the security zone.

```
[edit security]
lsdesignadmin1@host:ls-product-design# set zones security-zone ls-product-design-trust
application-tracking
```

3. Generate update messages at session start and at 5-minute intervals.

```
[edit security]
lsdesignadmin1@host:ls-product-design# set application-tracking first-update
```

Results

From configuration mode, confirm your configuration by entering the `show security` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this `show` command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
lsdesignadmin1@host:ls-product-design# show security
...
    application-tracking {
        first-update;
    }
...
    zones {
        security-zone ls-product-design-trust {
            ...
            application-tracking;
        }
    }
```

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

Verification

IN THIS SECTION

- Verifying AppTrack Operation | 358
- Verifying Security Flow Session Statistics | 358
- Verifying Application System Cache Statistics | 359
- Verifying the Status of Application Identification Counter Values | 359

To confirm that the configuration is working properly, perform these tasks:

Verifying AppTrack Operation

Purpose

View the AppTrack counters periodically to monitor tracking.

Action

From operational mode, enter the `show application-tracking counters` command.

Verifying Security Flow Session Statistics

Purpose

Compare byte and packet counts in logged messages with the session statistics from the `show security flow session` command output.

Action

From operational mode, enter the `show security flow session` command.

Verifying Application System Cache Statistics

Purpose

Compare cache statistics such as IP address, port, protocol, and service for an application from the `show services application-identification application-system-cache` command output.

Action

From operational mode, enter the `show services application-identification application-system-cache` command.

Verifying the Status of Application Identification Counter Values

Purpose

Compare session statistics for application identification counter values from the `show services application-identification counter` command output.

Action

From operational mode, enter the `show services application-identification counter` command.

SEE ALSO

Understanding Logical Systems Application Tracking Services | 348

User Logical Systems Configuration Overview | 51

IPv6 for Logical Systems

IN THIS SECTION

- IPv6 Addresses in Logical Systems Overview | 360
- Understanding IPv6 Dual-Stack Lite in Logical Systems | 361

- Example: Configuring IPv6 for the Primary, Interconnect, and User Logical Systems (Primary Administrators Only) | **362**
- Example: Configuring IPv6 Zones for a User Logical Systems | **373**
- Example: Configuring IPv6 Security Policies for a User Logical Systems | **378**
- Example: Configuring IPv6 Dual-Stack Lite for a User Logical Systems | **383**

IPv6 builds upon the functionality of IPv4, providing improvements to IP addressing, configuration and maintenance, and security. IPv6 supports extensions for authentication and data integrity, which enhance privacy and security. IPv6 uses 128-bit addresses and supports a virtually unlimited number of devices—2 to the 128th power. For more information, see the following topics:

IPv6 Addresses in Logical Systems Overview

IP version 6 (IPv6) increases the size of an IP address from the 32 bits that compose an IPv4 address to 128 bits. Each extra bit given to an address doubles the size of its address space. IPv6 has a much larger address space than the soon-to-be exhausted IPv4 address space.

IPv6 addresses can be configured in logical systems for the following features:

- Interfaces
- Firewall authentication
- Flows
- Routing (BGP only)
- Zones and security policies
- Screen options
- Network Address Translation (except for interface NAT)
- Administrative operations such as SSH, HTTPS, and other utilities
- Chassis clusters

NOTE: An IPv6 session consumes twice the memory of an IPv4 session. Therefore the number of sessions available for IPv6 is half the reserved and maximum quotas configured for the flow session resource in a security profile. Use the vty command **show usp flow resource usage cp-session** to check flow session usage.

SEE ALSO

Understanding IPv6 Address Space, Addressing, Address Format, and Address Types

Example: Configuring IPv6 for the Primary, Interconnect, and User Logical Systems (Primary Administrators Only) | **362**

Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (IPv6) (Primary Administrators Only) | **456**

Understanding IPv6 Dual-Stack Lite in Logical Systems | **361**

Understanding IPv6 Dual-Stack Lite in Logical Systems

IPv6 dual-stack lite (DS-Lite) allows migration to an IPv6 access network without changing end-user software. IPv4 users can continue to access IPv4 internet content using their current hardware, while IPv6 users are able to access IPv6 content. A DS-Lite softwire initiator at the customer edge encapsulates IPv4 packets into IPv6 packets while a softwire concentrator decapsulates the IPv4-in-IPv6 packets and also performs IPv4 NAT translations.

A specific softwire concentrator and the set of softwire initiators that connect with that softwire concentrator can belong to only one logical system. The primary administrator configures the maximum and reserved numbers of softwire initiators that can be connected to a softwire concentrator in a logical system using the `dslite-softwire-initiator` *configuration statement* at the [`edit system security-profile resources`] hierarchy level. The default maximum value is the system maximum; the default reserved value is 0.

NOTE: The primary administrator can configure a security profile for the primary logical system that specifies the maximum and reserved numbers of softwire initiators that can connect to a softwire concentrator configured for the primary logical system. The number of softwire initiators configured in the primary logical system count toward the maximum number of softwire initiators available on the device.

The user logical system administrator can configure softwire concentrators for their user logical system and the primary administrator can configure softwire concentrators for the primary logical system at the [`edit security softwires`] hierarchy level. The primary administrator can also configure softwire concentrators for a user logical system at the [`edit logical-systems` *`logical-system`* `security softwires`] hierarchy level.

NOTE: The softwire concentrator IPv6 address can match an IPv6 address configured on either a physical interface or a loopback interface.

SEE ALSO

Example: Configuring IPv6 Dual-Stack Lite for a User Logical Systems | **383**

Understanding Logical Systems Security Profiles (Primary Administrators Only) | **71**

Understanding IPv6 Dual-Stack Lite

Example: Configuring IPv6 for the Primary, Interconnect, and User Logical Systems (Primary Administrators Only)

IN THIS SECTION

- Requirements | **362**
- Overview | **363**
- Configuration | **365**
- Verification | **372**

This topic covers configuration of IPv6 interfaces, static routes, and routing instances for the primary and interconnect logical systems. It also covers configuration of IPv6 logical tunnel interfaces for user logical systems.

Requirements

Before you begin:

- See "SRX Series Logical Systems Primary Administrator Configuration Tasks Overview" on page 24 to understand how and where this procedure fits in the overall primary administrator configuration process.
- See "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System" on page 57.
- See "Understanding the Interconnect Logical System and Logical Tunnel Interfaces" on page 11.

Overview

IN THIS SECTION

- Topology | 364

This scenario shows how to configure interfaces for the logical systems on the device, including an interconnect logical system.

- For the interconnect logical system, the example configures logical tunnel interfaces lt-0/0/0.0, lt-0/0/0.2, and lt-0/0/0.4. The example configures a routing instance called vr and assigns the interfaces to it.

  Because the interconnect logical system acts as a virtual switch, it is configured as a VPLS routing instance type. The interconnect logical system's lt-0/0/0 interfaces are configured with ethernet-vpls as the encapsulation type. The corresponding peer lt-0/0/0 interfaces in the primary and user logical systems are configured with Ethernet as the encapsulation type.

  - lt-0/0/0.0 connects to lt-0/0/0.1 on the root logical system.
  - lt-0/0/0.2 connects to lt-0/0/0.3 on the LSYS1 logical system.
  - lt-0/0/0.4 connects to lt-0/0/0.5 on the LSYS2 logical system.
- For the primary logical system, called root-logical-system, the example configures ge-5/0/0 and assigns it to the vr0 routing instance. The example configures lt-0/0/0.1 to connect to lt-0/0/0.0 on the interconnect logical system and assigns it to the vr0 routing instance. The example configures static routes to allow for communication with other logical systems and assigns them to the vr0 routing instance.
- For the LSYS1 logical system, the example configures lt-0/0/0.3 to connect to lt-0/0/0.2 on the interconnect logical system.

- For the LSYS2 logical system, the example configures lt-0/0/0.5 to connect to lt-0/0/0.4 on the interconnect logical system.

Figure 9 on page 365 shows the topology for this deployment including virtual routers and their interfaces for all IPv6 logical systems.

Topology

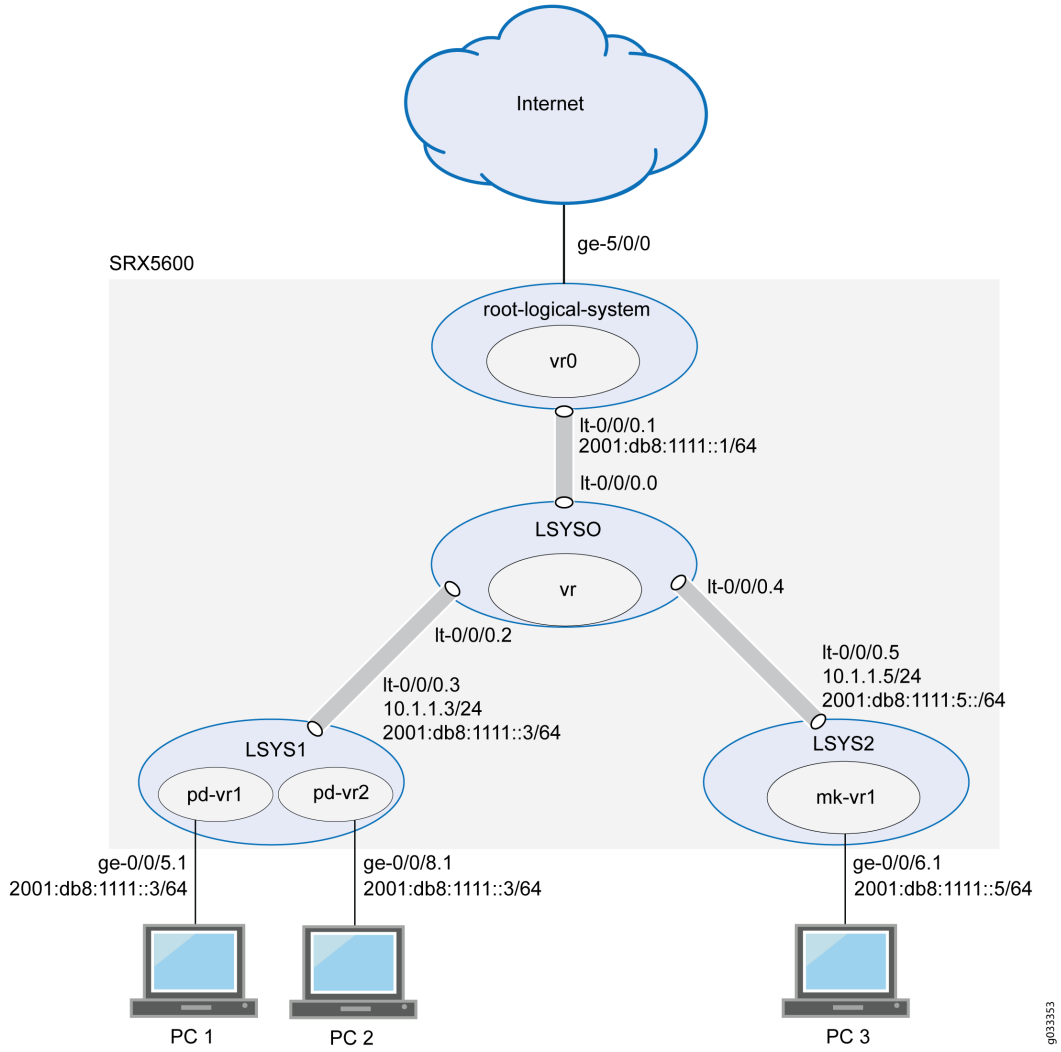

Figure 9: Configuring IPv6 Logical Tunnel Interfaces, Logical Interfaces, and Virtual Routers

Configuration

IN THIS SECTION

- Configuring Logical Tunnel Interfaces and a Routing Instance for the Interconnect Logical System | 365
- Configuring Interfaces, a Routing Instance, and Static Routes for the Primary Logical System | 367
- Configuring Logical Tunnel Interfaces for the User Logical Systems | 370

This topic explains how to configure interfaces for logical systems.

Configuring Logical Tunnel Interfaces and a Routing Instance for the Interconnect Logical System

CLI Quick Configuration

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

```
set forwarding-options family inet6 mode flow-based
set logical-systems LSYS0 interfaces lt-0/0/0 unit 0 encapsulation ethernet-vpls
set logical-systems LSYS0 interfaces lt-0/0/0 unit 0 peer-unit 1
set logical-systems LSYS0 interfaces lt-0/0/0 unit 2 encapsulation ethernet-vpls
set logical-systems LSYS0 interfaces lt-0/0/0 unit 2 peer-unit 3
set logical-systems LSYS0 interfaces lt-0/0/0 unit 4 encapsulation ethernet-vpls
set logical-systems LSYS0 interfaces lt-0/0/0 unit 4 peer-unit 5
set logical-systems LSYS0 routing-instances vr instance-type vpls
set logical-systems LSYS0 routing-instances vr interface lt-0/0/0.0
set logical-systems LSYS0 routing-instances vr interface lt-0/0/0.2
set logical-systems LSYS0 routing-instances vr interface lt-0/0/0.4
```

Step-by-Step Procedure

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

To configure the interconnect system lt-0/0/0 interfaces and routing instances:

1. Enable flow-based forwarding for IPv6 traffic.

```
[edit security]
user@host# set forwarding-options family inet6 mode flow-based
```

2. Configure the lt-0/0/0 interfaces.

```
[edit logical-systems LSYS0 interfaces]
user@host# set lt-0/0/0 unit 0 encapsulation ethernet-vpls
user@host# set lt-0/0/0 unit 0 peer-unit 1
user@host# set lt-0/0/0 unit 2 encapsulation ethernet-vpls
user@host# set lt-0/0/0 unit 2 peer-unit 3
user@host# set lt-0/0/0 unit 4 encapsulation ethernet-vpls
user@host# set lt-0/0/0 unit 4 peer-unit 5
```

3. Configure the routing instance for the interconnect logical system and add its lt-0/0/0 interfaces to it.

```
[edit logical-systems LSYS0 routing-instances]
user@host# set vr instance-type vpls
user@host# set vr interface lt-0/0/0.0
user@host# set vr interface lt-0/0/0.2
user@host# set vr interface lt-0/0/0.4
```

Results

From configuration mode, confirm your configuration by entering the `show logical-systems interconnect-logical-system` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

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

```
user@host# show logical-systems LSYS0
interfaces {
    lt-0/0/0 {
        unit 0 {
            encapsulation ethernet-vpls;
```

```
            peer-unit 1;
        }
        unit 2 {
            encapsulation ethernet-vpls;
            peer-unit 3;
            }
            unit 4 {
                encapsulation ethernet-vpls;
                peer-unit 5;
            }
        }
    }
        routing-instances {
            vr {
                instance-type vpls;
                interface lt-0/0/0.0;
                interface lt-0/0/0.2;
                interface lt-0/0/0.4;
            }
        }
```

Configuring Interfaces, a Routing Instance, and Static Routes for the Primary Logical System

CLI Quick Configuration

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

```
set interfaces ge-5/0/0 vlan-tagging
set interfaces ge-5/0/0 unit 0 vlan-id 600
set interfaces lt-0/0/0 unit 1 encapsulation Ethernet
set interfaces lt-0/0/0 unit 1 peer-unit 0
set interfaces lt-0/0/0 unit 1 family inet address 10.1.1.1/24
set interfaces lt-0/0/0 unit 1 family inet6 address 2001:db8:1111::1/64
set interfaces ge-5/0/0 unit 0 family inet address 10.99.99.1/24
set interfaces ge-5/0/0 unit 0 family inet6 address 2001:db8:9999::1/64
set routing-instances vr0 instance-type virtual-router
set routing-instances vr0 interface lt-0/0/0.1
set routing-instances vr0 interface ge-5/0/0.0
set routing-instances vr0 routing-options rib vr0.inet6.0 static route 2001:db8:777::/64 next-hop
```

```
2001:db8:1111::3
set routing-instances vr0 routing-options rib vr0.inet6.0 static route 2001:db8:888::/64 next-hop
2001:db8:1111::3
set routing-instances vr0 routing-options rib vr0.inet6.0 static route 2001:db8:666::/64 next-hop
2001:db8:1111::5
set routing-instances vr0 routing-options static route 192.168.7.0/24 next-hop 10.1.1.3
set routing-instances vr0 routing-options static route 192.168.8.0/24 next-hop 10.1.1.3
set routing-instances vr0 routing-options static route 192.168.6.0/24 next-hop 10.1.1.5
```

Step-by-Step Procedure

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

To configure the primary logical system interfaces:

1. Configure the primary (root) logical system and lt-0/0/0.1 interfaces.

```
[edit interfaces]
user@host# set ge-5/0/0 vlan-tagging
user@host# set ge-5/0/0 unit 0 vlan-id 600
user@host# set lt-0/0/0 unit 1 encapsulation Ethernet
user@host# set lt-0/0/0 unit 1 peer-unit 0
user@host# set lt-0/0/0 unit 1 family inet address 10.1.1.1/24
user@host# set lt-0/0/0 unit 1 family inet6 address 2001:db8:1111::1/64
user@host# set ge-5/0/0 unit 0 family inet address 10.99.99.1/24
user@host# set ge-5/0/0 unit 0 family inet6 address 2001:db8:9999::1/64
```

2. Configure a routing instance for the primary logical system, assign its interfaces to it, and configure static routes for it.

```
[edit interfaces routing-instances]
user@host# set vr0 instance-type virtual-router
user@host# set vr0 interface lt-0/0/0.1
user@host# set vr0 interface ge-5/0/0.0
user@host# set vr0 routing-options rib vr0.inet6.0 static route 2001:db8:1111:777/64 next-hop
2001:db8:1111::3
user@host# set vr0 routing-options rib vr0.inet6.0 static route 2001:db8:1111:888/64 next-hop
2001:db8:1111::3
user@host# set vr0 routing-options rib vr0.inet6.0 static route 2001:db8:1111:666/64 next-hop
2001:db8:1111::5
```

```
user@host# set vr0 routing-options static route 192.168.7.0/24 next-hop 10.1.1.3
user@host# set vr0 routing-options static route 192.168.8.0/24 next-hop 10.1.1.3
user@host# set vr0 routing-options static route 192.168.6.0/24 next-hop 10.1.1.5
```

Results

From configuration mode, confirm your configuration by entering the `show interfaces` and `show routing-instances` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
ge-5/0/0 {
    vlan-tagging;
    unit 0 {
        vlan-id 600;
        family inet {
            address 10.99.99.1/24;
        }
        family inet 6{
            address 2001:db8:9999:1/64;
        }
    }
}
lt-0/0/0 {
    unit 1 {
        encapsulation ethernet;
        peer-unit 0;
        family inet {
            address 10.1.1.1/24;
        }
        family inet 6{
            address 2001:db8:1111::1/64;
        }
    }
}
```

```
[edit]
user@host# show routing-instances
vr0 {
```

```
    instance-type virtual-router;
    interface ge-5/0/0.0;
    interface lt-0/0/0;
    routing-options {
        rib vr0.inet6.0 {
        static {
            route 2001:db8:1111:888/64 next-hop 2001:db8:1111:3;
            route 2001:db8:1111:777/64 next-hop 2001:db8:1111:3;
            route 2001:db8:1111:666/64 next-hop 2001:db8:1111:5;
        }
        }
        static {
            route 192.168.7.0/24 next-hop 10.1.1.3;
            route 192.168.8.0/24 next-hop 10.1.1.3;
            route 192.168.6.0/24 next-hop 10.1.1.5;
        }
    }
}
```

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

Configuring Logical Tunnel Interfaces for the User Logical Systems

CLI Quick Configuration

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

```
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 encapsulation ethernet
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 peer-unit 2
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 family inet address 10.1.1.3/24
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 family inet6 address 2001:db8:1111:2/64
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 encapsulation ethernet
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 peer-unit 4
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 family inet address 10.1.1.5/24
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 family inet6 address 2001:db8:1111::5/64
```

Step-by-Step Procedure

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

1. Configure the lt-0/0/0 interface for the first user logical system:

```
[edit logical-systems LSYS1 interfaces lt-0/0/0 unit 3]
user@host# set encapsulation ethernet
user@host# set peer-unit 2
user@host# set family inet address 10.1.1.3/24
user@host# set family inet6 address 2001:db8:1111:3::1/64
```

2. Configure the lt-0/0/0 interface for the second user logical system.

```
[edit logical-systems LSYS2 interfaces lt-0/0/0 unit 5]
user@host# set encapsulation ethernet
user@host# set peer-unit 4
user@host# set family inet address 10.1.1.5/24
user@host# set family inet6 address 2001:db8:1111::5/64
```

Results

From configuration mode, confirm your configuration by entering the `show logical-systems LSYS1 interfaces lt-0/0/0`, and `show logical-systems LSYS2 interfaces lt-0/0/0` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host#  show logical-systems LSYS1 interfaces lt-0/0/0
```

```
lt-0/0/0 {
    unit 3 {
        encapsulation ethernet;
        peer-unit 2;
        family inet {
            address 10.1.1.3/24;
        }
        family inet 6{
            address 2001:db8:1111::3/64;
        }
```

```
    }
}
```

```
user@host#  show logical-systems LSYS2 interfaces lt-0/0/0
```

```
lt-0/0/0 {
    unit 5 {
        encapsulation ethernet;
        peer-unit 4;
        family inet {
            address 10.1.1.5/24;
        }
        family inet 6{
            address 2001:db8:1111::5/64;
        }
    }
}
```

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

Verification

IN THIS SECTION

- Verifying That the Static Routes Configured for the Primary Administrator Are Correct | 372

Verifying That the Static Routes Configured for the Primary Administrator Are Correct

Purpose

Confirm that the configuration is working properly. Verify if you can send data from the primary logical system to the other logical systems.

Action

From operational mode, use the `ping` command.

SEE ALSO

Understanding the Primary Logical Systems and the Primary Administrator Role | 23
Understanding User Logical Systems and the User Logical System Administrator Role | 53
Understanding the Interconnect Logical System and Logical Tunnel Interfaces | 11
Example: Configuring IPv6 Zones for a User Logical Systems | 373
Example: Configuring IPv6 Security Policies for a User Logical Systems | 378

Example: Configuring IPv6 Zones for a User Logical Systems

IN THIS SECTION

Requirements | 373
Overview | 373
Configuration | 374

This example shows how to configure IPv6 zones for a user logical system.

Requirements

Before you begin:

- Log in to the user logical system as the user logical system administrator.

  See "User Logical Systems Configuration Overview" on page 51.

- Ensure that forwarding options for inet6 is flow-based. Otherwise, you must configure it and reset the device.

  Use the `show security forwarding-options` command to check the configuration.

  NOTE: Only the user logical system administrator can configure the forwarding options.

Overview

This example configures the ls-product-design user logical system described in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System" on page 57

This example creates the IPv6 zones and address books described in Table 24 on page 374.

Table 24: User Logical System Zone and Address Book Configuration

Feature	Name	Configuration Parameters
Zones	ls-product-design-trust	• Bind to interface ge-0/0/5.1. • TCP reset enabled.
	ls-product-design-untrust	• Bind to interface lt-0/0/0.3.
Address books	product-design-internal	• Address product-designers: 3002::1/96 • Attach to zone ls-product-design-trust
	product-design-external	• Address marketing: 3003::1/24 • Address accounting: 3004::1/24 • Address others: 3002::2/24 • Address set otherlsys: marketing, accounting • Attach to zone ls-product-design-untrust

Configuration

IN THIS SECTION

- Procedure | 375

Procedure

CLI Quick Configuration

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

```
set logical-system lsys1 security address-book product-design-internal address product-designers
3002::1/96
set logical-system lsys1 security address-book product-design-internal attach zone ls-product-
design-trust
set logical-system lsys1 security address-book product-design-external address marketing
3003::1/24
set logical-system lsys1 security address-book product-design-external address accounting
3004::1/24
set logical-system lsys1 security address-book product-design-external address others 3002::2/24
set logical-system lsys1 security address-book product-design-external address-set otherlsys
address marketing
set logical-system lsys1 security address-book product-design-external address-set otherlsys
address accounting
set logical-system lsys1 security address-book product-design-external attach zone ls-product-
design-untrust
set logical-system lsys1 security zones security-zone ls-product-design-trust tcp-rst
set logical-system lsys1 security zones security-zone ls-product-design-trust interfaces
ge-0/0/5.1
set logical-system lsys1 security zones security-zone ls-product-design-untrust interfaces
lt-0/0/0.3
```

Step-by-Step Procedure

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

To configure IPv6 zones in a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

```
lsdesignadmin1@host:ls-product-design> configure
lsdesignadmin1@host:ls-product-design#
```

2. Configure a security zone and assign it to an interface.

```
[edit logical-system lsys1 security zones]
lsdesignadmin1@host:ls-product-design# set security-zone ls-product-design-trust interfaces
ge-0/0/5.1
```

3. Configure the TCP-Reset parameter for the zone.

```
[edit logical-system lsys1 security zones security-zone ls-product-design-trust]
lsdesignadmin1@host:ls-product-design#  set tcp-rst
```

4. Configure a security zone and assign it to an interface.

```
[edit logical-system lsys1 security zones]
lsdesignadmin1@host:ls-product-design# set security-zone ls-product-design-untrust interfaces
lt-0/0/0.3
```

5. Create global address book entries.

```
[edit logical-system lsys1 security]
lsdesignadmin1@host:ls-product-design# set address-book product-design-internal address
product-designers 3002::1/96
lsdesignadmin1@host:ls-product-design# set address-book product-design-external address
marketing 3003::1/24
lsdesignadmin1@host:ls-product-design# set address-book product-design-external address
accounting 3004::1/24
lsdesignadmin1@host:ls-product-design# set address-book product-design-external address
others 3002::2/24
lsdesignadmin1@host:ls-product-design# set address-book product-design-external address-set
otherlsys address marketing
lsdesignadmin1@host:ls-product-design# set address-book product-design-external address-set
otherlsys address accounting
```

6. Attach address books to zones.

```
[edit logical-system lsys1 security]
lsdesignadmin1@host:ls-product-design#set address-book product-design-internal attach zone ls-product-design-trust
lsdesignadmin1@host:ls-product-design#set address-book product-design-external attach zone ls-product-design-untrust
```

Results

From configuration mode, confirm your configuration by entering the `show security zones` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
lsdesignadmin1@host:ls-product-design# show security zones
address-book {
    product-design-internal {
        address product-designers 3002::1/96;
        attach {
            zone ls-product-design-trust;
        }
    }
    product-design-external {
        address marketing 3003::1/24;
        address accounting 3004::1/24;
        address others 3002::2/24;
        address-set otherlsys {
            address marketing;
            address accounting;
        }
        attach {
            zone ls-product-design-untrust;
        }
    }
}
zones {
    security-zone ls-product-design-trust {
        tcp-rst;
        interfaces {
            ge-0/0/5.1;
        }
```

```
    }
    security-zone ls-product-design-untrust {
        interfaces {
            lt-0/0/0.3;
        }
    }
}
```

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

SEE ALSO

- Understanding Logical Systems Zones | 149
- User Logical Systems Configuration Overview | 51
- Example: Configuring IPv6 for the Primary, Interconnect, and User Logical Systems (Primary Administrators Only) | 362
- Example: Configuring IPv6 Security Policies for a User Logical Systems | 378

Example: Configuring IPv6 Security Policies for a User Logical Systems

IN THIS SECTION

- Requirements | 378
- Overview | 379
- Configuration | 380
- Verification | 382

This example shows how to configure IPv6 security policies for a user logical system.

Requirements

Before you begin:

- Log in to the user logical system as the logical system administrator.

  See "User Logical Systems Configuration Overview" on page 51.

- Use the `show system security-profiles policy` command to see the security policy resources allocated to the logical system.
- Configure zones and address books.

  See "Example: Configuring IPv6 Zones for a User Logical Systems" on page 373

Overview

This example shows how to configure the security policies described in Table 25 on page 379.

Table 25: User Logical System Security Policies Configuration

Policy Name	Configuration Parameters
permit-all-to-otherlsys	Permit the following traffic: • From zone: ls-product-design-trust • To zone: ls-product-design-untrust • Source address: product-designers • Destination address: otherlsys • Application: any
permit-all-from-otherlsys	Permit the following traffic: • From zone: ls-product-design-untrust • To zone: ls-product-design-trust • Source address: otherlsys • Destination address: product-designers • Application: any

Configuration

IN THIS SECTION

- Procedure | 380

Procedure

CLI Quick Configuration

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

```
set logical-systems lsys1 security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust policy permit-all-to-otherlsys match source-address product-designers
set logical-systems lsys1 security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust policy permit-all-to-otherlsys match destination-address otherlsys
set logical-systems lsys1 security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust policy permit-all-to-otherlsys match application any
set logical-systems lsys1 security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust policy permit-all-to-otherlsys then permit
set logical-systems lsys1 security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust policy permit-all-from-otherlsys match source-address otherlsys
set logical-systems lsys1 security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust policy permit-all-from-otherlsys match destination-address product-designers
set logical-systems lsys1 security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust policy permit-all-from-otherlsys match application any
set logical-systems lsys1 security policies from-zone ls-product-design-untrust to-zone ls-product-d esign-trust policy permit-all-from-otherlsys then permit
```

Step-by-Step Procedure

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

To configure IPv6 security policies for a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

```
lsdesignadmin1@host:ls-product-design> configure
lsdesignadmin1@host:ls-product-design#
```

2. Configure a security policy that permits traffic from the ls-product-design-trust zone to the ls-product-design-untrust zone.

```
[edit logical-systems lsys1 security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust]
lsdesignadmin1@host:ls-product-design# set policy permit-all-to-otherlsys match source-address product-designers
lsdesignadmin1@host:ls-product-design# set policy permit-all-to-otherlsys match destination-address otherlsys
lsdesignadmin1@host:ls-product-design# set policy permit-all-to-otherlsys match application any
lsdesignadmin1@host:ls-product-design# set policy permit-all-to-otherlsys then permit
```

3. Configure a security policy that permits traffic from the ls-product-design-untrust zone to the ls-product-design-trust zone.

```
[edit logical-systems lsys1 security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust]
lsdesignadmin1@host:ls-product-design# set policy permit-all-from-otherlsys match source-address otherlsys
lsdesignadmin1@host:ls-product-design# set policy permit-all-from-otherlsys match destination-address product-designers
lsdesignadmin1@host:ls-product-design# set policy permit-all-from-otherlsys match application any
lsdesignadmin1@host:ls-product-design# set policy permit-all-from-otherlsys then permit
```

Results

From configuration mode, confirm your configuration by entering the `show security policies` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
lsdesignadmin1@host:ls-product-design# show security policies
from-zone ls-product-design-trust to-zone ls-product-design-untrust {
    policy permit-all-to-otherlsys {
        match {
            source-address product-designers;
            destination-address otherlsys;
            application any;
        }
        then {
            permit;
        }
    }
}
from-zone ls-product-design-untrust to-zone ls-product-design-trust {
    policy permit-all-from-otherlsys {
        match {
            source-address otherlsys;
            destination-address product-designers;
            application any;
        }
        then {
            permit;
        }
    }
}
```

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

Verification

IN THIS SECTION

- Verifying Policy Configuration | 383

Verifying Policy Configuration

Purpose

Verify information about policies and rules.

Action

From operational mode, enter the `show security policies detail` command to display a summary of all policies configured on the logical system.

SEE ALSO

- Understanding Logical Systems Security Policies | 213
- User Logical Systems Configuration Overview | 51
- Troubleshooting Security Policies
- Example: Configuring IPv6 Zones for a User Logical Systems | 373
- Example: Configuring IPv6 for the Primary, Interconnect, and User Logical Systems (Primary Administrators Only) | 362

Example: Configuring IPv6 Dual-Stack Lite for a User Logical Systems

IN THIS SECTION

- Requirements | 383
- Overview | 384
- Configuration | 384
- Verification | 385

This example shows how to configure a softwire concentrator for a user logical system.

Requirements

Before you begin:

- Log in to the user logical system as the user logical system administrator. See "User Logical Systems Configuration Overview" on page 51.
- Use the `show system security-profile dslite-softwire-initiator` command to see the number softwire initiators that can be connected to a softwire concentrator in the logical system.

Overview

This example shows how to configure a softwire concentrator to decapsulate IPv4-in-IPv6 packets in the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System" on page 57. The IPv6 address of the softwire concentrator is 3000::1 and the name of the softwire configuration is sc_1.

Configuration

IN THIS SECTION

- Procedure | 384

Procedure

CLI Quick Configuration

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

```
set security softwires softwire-name sc_1 softwire-concentrator 3000::1  softwire-type IPv4-in-IPv6
```

Step-by-Step Procedure

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

To configure an IPv6 DS-Lite softwire concentrator:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

```
lsdesignadmin1@host:ls-product-design> configure
lsdesignadmin1@host:ls-product-design#
```

2. Specify the address of the softwire concentrator and the softwire type.

```
[edit security]
lsdesignadmin1@host:ls-product-design# set softwires softwire-name sc_1 softwire-concentrator
3000::1 softwire-type IPv4-in-IPv6
```

Results

From configuration mode, confirm your configuration by entering the `show security softwires` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
lsdesignadmin1@host:ls-product-design# show security softwires
softwire-name sc_1 {
    softwire-concentrator 3000::1;
    softwire-type IPv4-in-IPv6;
}
```

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

Verification

IN THIS SECTION

- Verifying the DS-Lite Configuration | 386

Verifying the DS-Lite Configuration

Purpose

Verify that the softwire initiators can connect to the softwire concentrator configured in the user logical system.

Action

From operational mode, enter the `show security softwires` command.

If a softwire initiator is not connected, the operational output looks like this:

```
lsdesignadmin1@host:ls-product-design> show security softwires
Softwire Name            SC Address         Status   Number of SI connected
  sc_1                      3000::1         Active         0
```

If a softwire initiator is connected, the operational output looks like this:

```
lsdesignadmin1@host:ls-product-design> show security softwires
Softwire Name            SC Address         Status   Number of SI connected
  sc_1                      3000::1         Connected      1
```

SEE ALSO

Understanding IPv6 Dual-Stack Lite in Logical Systems | 361

User Logical Systems Configuration Overview | 51

RELATED DOCUMENTATION

Understanding Logical Systems Zones | 149

SSL Proxy for Logical Systems

IN THIS SECTION

- Understanding SSL Forward and Reverse Proxy for Logical Systems | 387
- Example: Configuring SSL Forward and Reverse Proxy for Logical Systems | 388

Secure Sockets Layer (SSL) is an application-level protocol that provides encryption technology for the Internet. For more information, see the following topics:

Understanding SSL Forward and Reverse Proxy for Logical Systems

SSL proxy acts as an intermediary, performing SSL encryption and decryption between the client and the server. SSL, also called Transport Layer Security (TLS), ensures the secure transmission of data between a client and a server through a combination of privacy, authentication, confidentiality, and data integrity.

SSL proxy is a transparent proxy that performs SSL encryption and decryption between the client and the server as follows:

- Reverse proxy is an inbound session, that is, externally initiated SSL sessions from the Internet to the local server.

  The proxy model implementation for server protection (often called reverse proxy) is supported on SRX Series devices to provide improved handshaking and support for more protocol versions.

- Forward proxy is an outbound session, that is, locally initiated SSL session to the Internet.

SSL proxy works transparently between the client and the server. All requests from a client first go to the proxy server; the proxy server evaluates the request, and if the request is valid, forwards the request to the outbound side. Similarly, inbound requests are also evaluated by the proxy server. Both client and server interpret that they are communicating with each other; however, it is the SSL proxy that functions between the two.

Example: Configuring SSL Forward and Reverse Proxy for Logical Systems

IN THIS SECTION

- Requirements | 388
- Overview | 388
- Configuration | 388
- Verification | 392

This example shows how to configure SSL proxy to enable server protection. A reverse proxy protects servers by hiding the details of the servers from the clients, there by adding an extra layer of security and the purpose of a forward proxy is to manage traffic to the client systems.

Requirements

To configure an SSL reverse and forward proxy, you must:

- Load the server certificates and their keys into SRX Series device's certificate repository.
- Attach the server certificate identifiers to the SSL proxy profile.
- Apply SSL proxy profile as application services in a security policy.

Overview

This example shows how to configure reverse proxy to enable server protection and forward proxy is for client protection. It shows how to configure an SSL proxy profile and apply it at the security policy rule level. For server protection, additionally, server certificates with private keys must be configured.

Configuration

IN THIS SECTION

- CLI Quick Configuration | 389
- Configuring the SSL Reverse and Forward Proxy | 390
- Results | 391

CLI Quick Configuration

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

```
set logical-systems LSYS1 services ssl proxy profile ssl-fp-profile root-ca new-srvr-cert
set logical-systems LSYS1 services ssl proxy profile ssl-fp-profile actions ignore-server-auth-
failure
set logical-systems LSYS1 services ssl proxy profile ssl-rp-profile actions log all
set logical-systems LSYS1 security log mode event
set logical-systems LSYS1 security policies from-zone trust to-zone untrust policy 1 match
source-address any
set logical-systems LSYS1 security policies from-zone trust to-zone untrust policy 1 match
destination-address any
set logical-systems LSYS1 security policies from-zone trust to-zone untrust policy 1 match
application any
set logical-systems LSYS1 security policies from-zone trust to-zone untrust policy 1 then permit
application-services ssl-proxy profile-name ssl-rp-profile
set logical-systems LSYS1 security policies from-zone trust to-zone untrust policy 1 then log
session-init
set logical-systems LSYS1 security policies from-zone trust to-zone untrust policy 1 then log
session-close
set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy 1 match
source-address any
set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy 1 match
destination-address any
set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy 1 match
application any
set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy 1 then permit
application-services ssl-proxy profile-name ssl-rp-profile
set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy 1 then log
session-init
set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy 1 then log
session-close
```

Configuring the SSL Reverse and Forward Proxy

Step-by-Step Procedure

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

To configure the SSL Proxy:

1. Configure the SSL Reverse Proxy.

```
[edit logical-systems LSYS1]
user@host# set logical-systems LSYS1 services ssl proxy profile ssl-rp-profile  actions log
all
user@host# set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy
1 match source-address any
user@host# set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy
1 match destination-address any
user@host# set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy
1 match application any
user@host# set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy
1 then permit application-services ssl-proxy profile-name ssl-rp-profile
```

2. Configure the SSL Forward Proxy.

```
[edit logical-systems LSYS1]
user@host# set logical-systems LSYS1 services ssl proxy profile ssl-fp-profile  root-ca new-
srvr-cert
user@host# set logical-systems LSYS1 services ssl proxy profile ssl-fp-profile  actions
ignore-server-auth-failure
user@host# set logical-systems LSYS1 security log mode event
user@host# set logical-systems LSYS1 security policies from-zone trust to-zone untrust policy
1 match source-address any
user@host# set logical-systems LSYS1 security policies from-zone trust to-zone untrust policy
1 match destination-address any
user@host# set logical-systems LSYS1 security policies from-zone trust to-zone untrust policy
1 match application any
user@host# set logical-systems LSYS1 security policies from-zone trust to-zone untrust policy
1 then permit application-services idp
user@host# set logical-systems LSYS1 security policies from-zone trust to-zone untrust policy
1 then permit application-services ssl-proxy profile-name ssl-rp-profile
```

```
user@host# set logical-systems LSYS1 security policies from-zone trust to-zone untrust policy
1 then log session-init
user@host# set logical-systems LSYS1 security policies from-zone trust to-zone untrust policy
1 then log session-close
user@host# set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy
1 then log session-init
user@host# set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy
1 then log session-close
```

Results

From configuration mode, confirm your configuration by entering the `show logical-system LSYS1 services ssl proxy` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

You must configure either `root-ca` (forward proxy) or `server-certificate` (reverse proxy) in an SSL proxy profile. Otherwise, the commit check fails.

```
user@host# show logical-systems LSYS1 services ssl proxy
profile ssl-rp-profile {
    server-certificate ssl-inspect-sp1;  { # For reverse proxy. No root-ca is needed.
    actions {
        log {
            all;
        }
    }
}
profile ssl-fp-profile { # For forward proxy. No server cert/key is needed.
    root-ca new-srvr-cert;
    actions {
        ignore-server-auth-failure;
        log {
            all;
        }
    }
}
```

Verification

IN THIS SECTION

- Verifying the SSL Proxy Configuration on the Device | **392**

Verifying the SSL Proxy Configuration on the Device

Purpose

Viewing the SSL reverse proxy statistics on the SRX Series device.

Action

You can view the SSL proxy statistics by using the `show services ssl proxy statistics logical-system` command.

```
user@host> show services ssl proxy statistics logical-system LSYS1
PIC:spu-3 fpc[0] pic[3] ------
        sessions matched                                1
        sessions bypassed:non-ssl                       0
        sessions bypassed:mem overflow                  0
        sessions bypassed:low memory                    0
        sessions created                                1
        sessions ignored                                0
        sessions active                                 1
        sessions dropped                                0
        sessions whitelisted                            0
        whitelisted url category match                  0
        default profile hit                             0
        session dropped no default profile              0
        policy hit no profile configured                0
```

SEE ALSO

Example: Loading CA and Local Certificates Manually

Example: Configuring a Device for Peer Certificate Chain Validation

ICAP Redirects for Logical Systems

IN THIS SECTION

- ICAP Redirect Support for Logical Systems | 393
- Example: Configuring ICAP Redirect Service on SRX Devices | 395

ICAP is a lightweight protocol used to extend transparent proxy servers, thereby freeing up resources. For more information, see the following topics:

ICAP Redirect Support for Logical Systems

IN THIS SECTION

- Limitations of SSL Proxy with Logical Systems | 394

Starting in Junos OS Release 18.3R1, SRX Series devices support the Internet Content Adaptation Protocol (ICAP) service redirect when the device is configured for logical systems.

ICAP redirect profile is only allowed to attach on the policy which belongs to the same logical system. This profile is applied to a security policy as application services for the permitted traffic. The ICAP profile defines the settings that allow the ICAP server to process request messages, response messages, fallback options (in case of a timeout), connectivity issues, too many requests, or any other conditions.

Secure Sockets Layer (SSL) is an application-level protocol that provides encryption technology for the Internet. SSL proxy acts as an intermediary, performing SSL encryption and decryption between the client and the server. SSL, also called Transport Layer Security (TLS), ensures the secure transmission of

data between a client and a server through a combination of privacy, authentication, confidentiality, and data integrity. SSL relies on certificates and private-public key exchange pairs for this level of security. SSL proxy acts as an intermediary, performing SSL encryption and decryption between the client and the server, but neither the server nor the client can detect its presence.

ICAP redirect services has the dependency on SSL proxy to build secure connections. Because the SSL proxy is not supported on user logical systems in Junos OS Release 18.3R1, ICAP redirect works with clear text connections or with shared certificates in Junos OS Release 18.3R1.

The following sequences are involved in a typical ICAP redirect scenario:

1. The user opens a connection to a Website on the internet.
2. The request goes through the SRX Series device that is acting as a proxy server.
3. The SRX Series device receives information from the end-host, encapsulates the message and forwards the encapsulated ICAP message to the third-party on-premise ICAP server.
4. The ICAP server receives the ICAP request and analyzes it.
5. If the request does not contain any confidential information, the ICAP server sends it back to the proxy server, and directs the proxy server to send the HTTP to the internet.
6. If the request contains confidential information, you can choose to take action (block, permit and log) as per your requirement.

Limitations of SSL Proxy with Logical Systems

Following are the limitations for using ICAP redirect service for user logical systems:

- SSL Proxy is supported only on primary logical system in Junos OS Release 18.3R1.
- SSL profile configured to provide a secure connection to the ICAP server is not supported on user logical systems in Junos OS Release 18.3R1.

SEE ALSO

SSL Proxy

Example: Configuring ICAP Redirect Service on SRX Devices

IN THIS SECTION

- Requirements | 395
- Overview | 395
- Configuration | 396
- Verification | 400

This example shows how to define an ICAP redirect profile for an SRX Series device.

Requirements

This example uses the following hardware and software components:

- SRX Series device with Junos OS Release 18.3R1 or later. This configuration example is tested for Junos OS Release 18.3R1.

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you configure an ICAP redirect profile in logical systems and apply these profiles as application services in the security policy for the permitted traffic.

Table 26 on page 395 lists the details of the parameters used in this example.

Table 26: ICAP Redirect Configuration Parameters

Parameters	Names	Description
Profile	icap-pf1	The ICAP server profile allows the ICAP server to process request messages, response messages, fallback options and so on, for the permitted traffic. This profile is applied as an application service in the security policy.
Server name	icap-svr1 icap-svr2	The machine name of the remote ICAP host. Client's request is redirected to this ICAP server.

Table 26: ICAP Redirect Configuration Parameters *(Continued)*

Parameters	Names	Description
Server IP address	192.0.2.2/24 192.0.2.179/24	The IP address of the remote ICAP host. Client's request is redirected to this ICAP server.
Logical system name	LSYS1	Displays the logical system name which belongs to the same profile.
Security policy	sp1	In a security policy, apply the SSL proxy profile and ICAP redirect profile. to the permitted traffic.

Configuration

IN THIS SECTION

- Procedure | 396

Procedure

CLI Quick Configuration

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

```
set logical-systems LSYS1 services icap-redirect profile icap-pf1 server icap-svr1 host
192.0.2.2/24
set logical-systems LSYS1 services icap-redirect profile icap-pf1 server icap-svr1 reqmod-uri
echo
set logical-systems LSYS1 services icap-redirect profile icap-pf1 server icap-svr1 respmod-uri
echo
set logical-systems LSYS1 services icap-redirect profile icap-pf1 server icap-svr1 sockets 64
set logical-systems LSYS1 services icap-redirect profile icap-pf1 server icap-svr2 host
192.0.2.179/24
set logical-systems LSYS1 services icap-redirect profile icap-pf1 server icap-svr2 reqmod-uri
```

```
echo
set logical-systems LSYS1 services icap-redirect profile icap-pf1 server icap-svr2 respmod-uri
echo
set logical-systems LSYS1 services icap-redirect profile icap-pf1 server icap-svr2 sockets 64
set logical-systems LSYS1 services icap-redirect profile icap-pf1 server icap-svr2 tls-profile
dlp_ssl
set logical-systems LSYS1 services icap-redirect profile icap-pf1 http redirect-request
set logical-systems LSYS1 services icap-redirect profile icap-pf1 http redirect-response
set logical-systems LSYS1 security policies from-zone trust to-zone untrust policy sec_policy
match source-address any
set logical-systems LSYS1 security policies from-zone trust to-zone untrust policy sec_policy
match destination-address any
set logical-systems LSYS1 security policies from-zone trust to-zone untrust policy sec_policy
match application any
set logical-systems LSYS1 security policies from-zone trust to-zone untrust policy sec_policy
then permit application-services ssl-proxy profile-name ssl-inspect-profile
set logical-systems LSYS1 security policies from-zone trust to-zone untrust policy sec_policy
then permit application-services icap-redirect icap-pf1
set logical-systems LSYS1 security policies default-policy permit-all
set logical-systems LSYS1 security zones security-zone trust host-inbound-traffic system-
services all
set logical-systems LSYS1 security zones security-zone trust host-inbound-traffic protocols all
set logical-systems LSYS1 security zones security-zone trust interfaces xe-5/0/0.0
set logical-systems LSYS1 security zones security-zone untrust host-inbound-traffic system-
services all
set logical-systems LSYS1 security zones security-zone untrust host-inbound-traffic protocols all
set logical-systems LSYS1 security zones security-zone untrust interfaces xe-5/0/1.0
set logical-systems LSYS1 interfaces xe-5/0/0 unit 0 family inet address 192.0.2.1/8
set logical-systems LSYS1 interfaces xe-5/0/0 unit 0 family inet6 address 2001:db8::1/64
set logical-systems LSYS1 interfaces xe-5/0/1 unit 0 family inet address 198.51.100.1/8
set logical-systems LSYS1 interfaces xe-5/0/1 unit 0 family inet6 address 2001:db8::2/64
```

Step-by-Step Procedure

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

To configure the ICAP redirect service:

1. Configure the ICAP redirect profile for the first server (icap-svr1).

```
[edit logical-systems LSYS1 services]
user@host# set icap-redirect profile icap-pf1 server icap-svr1 host 192.0.2.2/24
user@host# set icap-redirect profile icap-pf1 server icap-svr1 reqmod-uri echo
user@host# set icap-redirect profile icap-pf1 server icap-svr1 respmod-uri echo
user@host# set icap-redirect profile icap-pf1 server icap-svr1 sockets 64
```

2. Configure the ICAP redirect profile for the second server (icap-svr2).

```
[edit logical-systems LSYS1 services]
user@host# set icap-redirect profile icap-pf1 server icap-svr2 host 192.0.2.179/24
user@host# set icap-redirect profile icap-pf1 server icap-svr2 reqmod-uri echo
user@host# set icap-redirect profile icap-pf1 server icap-svr2 respmod-uri echo
user@host# set icap-redirect profile icap-pf1 server icap-svr2 sockets 64
user@host# set icap-redirect profile icap-pf1 server icap-svr2 tls-profile dlp_ssl
```

3. Configure the redirect request and the redirect response for the HTTP traffic.

```
[edit logical-systems LSYS1 services]
user@host# set icap-redirect profile icap-pf1 http redirect-request
user@host# set icap-redirect profile icap-pf1 http redirect-response
```

4. Configure a security policy to apply application services for the ICAP redirect to the permitted traffic.

```
[edit logical-systems LSYS1 security]
user@host# set policies from-zone trust to-zone untrust policy sec_policy match source-
address any
user@host# set policies from-zone trust to-zone untrust policy sec_policy match destination-
address any
user@host# set policies from-zone trust to-zone untrust policy sec_policy match application
any
user@host# set policies from-zone trust to-zone untrust policy sec_policy then permit
application-services ssl-proxy profile-name ssl-inspect-profile
user@host# set policies from-zone trust to-zone untrust policy sec_policy then permit
application-services icap-redirect icap-pf1
user@host# set policies default-policy permit-all
```

5. Configure zones.

```
[edit logical-systems LSYS1 security]
user@host# set zones security-zone trust host-inbound-traffic system-services all
user@host# set zones security-zone trust host-inbound-traffic protocols all
user@host# set zones security-zone trust interfaces xe-5/0/0.0
user@host# set zones security-zone untrust host-inbound-traffic system-services all
user@host# set zones security-zone untrust host-inbound-traffic protocols all
user@host# set zones security-zone untrust interfaces xe-5/0/1.0
```

6. Configure interfaces.

```
[edit logical-systems LSYS1]
user@host# set interfaces xe-5/0/0 unit 0 family inet address 192.0.2.1/8
user@host# set interfaces xe-5/0/0 unit 0 family inet6 address 2001:db8::1/64
user@host# set interfaces xe-5/0/1 unit 0 family inet address 198.51.100.1/8
user@host# set interfaces xe-5/0/1 unit 0 family inet6 address 2001:db8::2/64
```

Results

From configuration mode, confirm your configuration by entering the `show logical-systems LSYS1 services icap-redirect`, `show logical-systems LSYS1 security policies`, `show logical-systems LSYS1 security zones`, and `show logical-systems LSYS1 interfaces` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show logical-systems LSYS1 services icap-redirect
profile icap-pf1 {
    server icap-svr1 {
        host 192.0.2.2/24;
        reqmod-uri echo;
        respmod-uri echo;
        sockets 64;
    }
    server icap-svr2 {
        host 192.0.2.179/24;
        reqmod-uri echo;
        respmod-uri echo;
        sockets 64;
        tls-profile dlp_ssl;
```

```
    }
    http {
        redirect-request;
        redirect-response;
    }
}
```

```
from-zone trust to-zone untrust {
    policy sec_policy {
        match {
            source-address any;
            destination-address any;
            application any;
        }
        then {
            permit {
                application-services {
                    ssl-proxy {
                        profile-name ssl-inspect-profile;
                    }
                    icap-redirect icap-pf1;
                }
            }
        }
    }
}
default-policy {
    permit-all;
}
```

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

Verification

IN THIS SECTION

- Verifying ICAP Redirect Configuration | **401**

Verifying ICAP Redirect Configuration

Purpose

Verify that the ICAP redirect service is configured on the device.

Action

From operational mode, enter the `show services icap-redirect status logical-system` and `show services icap-redirect statistic logical-system` commands.

```
user@host> show services icap-redirect status logical-system LSYS1

ICAP Status :
         spu-1 Profile: icap-pf1 Server: icap-svr1 : UP
ICAP Status :
         spu-2 Profile: icap-pf1 Server: icap-svr1 : UP
ICAP Status :
         spu-3 Profile: icap-pf1 Server: icap-svr1 : UP

user@host> show services icap-redirect statistic logical-system LSYS1

ICAP Redirect statistic:
  Message Redirected                 : 12
      Message REQMOD Redirected      : 6
      Message RESPMOD Redirected     : 6
  Message Received                   : 12
      Message REQMOD Received        : 6
      Message RESPMOD Received       : 6
Fallback:      permit           log-permit       reject
  Timeout         0                 0                0
  Connectivity    0                 0                0
  Default         0                 0                0
```

Meaning

The status `Up` indicates that the ICAP redirect service is enabled. The `Message Redirected` and the `Message Received` fields show the number of HTTP requests that have passed through the ICAP channel.

RELATED DOCUMENTATION

Example: Configuring Logical Systems Security Profiles (Primary Administrators Only) | 77

AppQoS for Logical Systems

IN THIS SECTION

- Application Quality of Service Support for Logical Systems Overview | 402
- Example: Configure Application Quality of Service for Logical Systems | 403

Application quality of service (AppQoS) enable you to identify and control access to specific applications and provides the granularity of the stateful firewall rule base to match and enforce quality of service (QoS) at the application layer. AppQoS feature expands the capability of Junos OS class of service (CoS) for logical systems.

Application Quality of Service Support for Logical Systems Overview

The application quality of service (AppQoS) feature expands the capability of Junos OS class of service (CoS) for logical systems. This includes marking DSCP values based on Layer-7 application types, honoring application-based traffic through loss priority settings, and controlling transfer rates on egress PICs based on Layer-7 application types.

When a network experiences congestion and delay, some packets must be dropped. Junos OS CoS allows you to divide traffic into classes and offer various levels of throughput and packet loss when congestion occurs. This allows packet loss to happen according to the rules you configure.

Logical system enables you to partition a single device into multiple domains to perform security and routing functions.

Starting in Junos OS Release 19.3R1, AppQoS is supported when the SRX Series device is configured with logical system. You can configure a default AppQoS rule set to manage the application- traffic-control within the logical system. AppQoS provides the ability to prioritize and meter the application traffic to provide better service to business-critical or high-priority application traffic.

AppQoS rule sets are included in the logical system to implement application-aware quality-of-service control. You can configure a rule set with rules under the application-traffic-control option, and attach the AppQoS rule set to a logical system as an application service. If the traffic matches the specified application the application-aware quality of service is applied for logical system.

For AppQoS, traffic is grouped based on rules that associate a defined forwarding class with selected applications for logical system. The match criteria for the rule includes one or more applications. When traffic from a matching application encounters the rule, the rule action sets the forwarding class, and remarks the DSCP value and loss priority to values appropriate for the application.

The AppQoS DSCP rewriter conveys a packet's quality of service through both the forwarding class and a loss priority. The AppQoS rate-limiting parameters control the transmission speed and volume for its associated queues for logical system. The default AppQoS rule set is leveraged from one of the existing AppQoS rule sets, which are configured under the `[edit class-of-service application-traffic-control]` hierarchy level.

Rate limiters are applied in rules based on the application of the traffic for logical system. Two rate limiters are applied for each session: `client-to-server` and `server-to-client`. This usage allows traffic in each direction to be provisioned separately.

Example: Configure Application Quality of Service for Logical Systems

IN THIS SECTION

- Requirements | **403**
- Overview | **404**
- Configuration | **404**
- Verification | **408**

This example shows how to enable application quality of service (AppQoS) within a logical system to provide prioritization and rate limiting for the traffic.

Requirements

This example uses the following hardware and software components:

- An SRX Series device configured with logical systems.
- Junos OS Release 19.3R1 and later releases.

Before you begin:

- Read the "Application Quality of Service Support for Logical Systems Overview" on page 402 to understand how and where this procedure fits in the overall support for AppQos.

  No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you configure an AppQoS rule set and invoke AppQoS as an application service in the logical system. You configure the class of service (CoS) for logical system. The AppQoS rule sets are included in the logical system to implement application-aware quality-of-service control.

Configuration

IN THIS SECTION

- CLI Quick Configuration | **404**
- Configuring AppQoS with a Logical System | **405**
- Results | **406**

CLI Quick Configuration

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

```
set logical-systems LSYS1 class-of-service application-traffic-control rate-limiters HTTP-BW-RL
bandwidth-limit 512
set logical-systems LSYS1 class-of-service application-traffic-control rule-sets RS1 rule RL1
match application junos:HTTP
set logical-systems LSYS1 class-of-service application-traffic-control rule-sets RS1 rule RL1
then forwarding-class best-effort
set logical-systems LSYS1 class-of-service application-traffic-control rule-sets RS1 rule RL1
then dscp-code-point 001000
set logical-systems LSYS1 class-of-service application-traffic-control rule-sets RS1 rule RL1
then loss-priority high
set logical-systems LSYS1 class-of-service application-traffic-control rule-sets RS1 rule RL1
then log
set logical-systems LSYS1 class-of-service application-traffic-control rule-sets RS1 rule RL1
```

```
then rate-limit server-to-client HTTP-BW-RL
set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy from_internet
match source-address any
set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy from_internet
match destination-address any
set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy from_internet
match application any
set logical-systems LSYS1 security policies from-zone trust to-zone trust policy p1 match
dynamic-application junos:web
set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy from_internet
then permit application-services application-traffic-control rule-set RS1
```

Configuring AppQoS with a Logical System

Step-by-Step Procedure

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

To configure AppQoS with a Logical System:

1. Configure the AppQoS real-time run information about application rate limiting of current or recent sessions for logical system LSYS1.

```
user@host# set logical-systems LSYS1 class-of-service application-traffic-control rate-
limiters HTTP-BW-RL bandwidth-limit 512
```

2. Configure the AppQoS rules and application match criteria for logical system LSYS1.

```
user@host# set logical-systems LSYS1 class-of-service application-traffic-control rule-sets
RS1 rule RL1 match application junos:HTTP
```

3. Configure the AppQoS rules and the forwarding class for logical system LSYS1.

```
user@host# set logical-systems LSYS1 class-of-service application-traffic-control rule-sets
RS1 rule RL1 then forwarding-class best-effort
```

4. Configure the AppQoS rules and the dscp-code-point for logical system LSYS1.

```
user@host# set logical-systems LSYS1 class-of-service application-traffic-control rule-sets
RS1 rule RL1 then dscp-code-point 001000
```

5. Configure the AppQoS rules and the loss priority for logical system LSYS1.

```
user@host# set logical-systems LSYS1 class-of-service application-traffic-control rule-sets
RS1 rule RL1 then loss-priority high
```

6. Assign the rate limiters for rule-sets.

```
user@host# set logical-systems LSYS1 class-of-service application-traffic-control rule-sets
RS1 rule RL1 then log
user@host# set logical-systems LSYS1 class-of-service application-traffic-control rule-sets
RS1 rule RL1 then rate-limit server-to-client HTTP-BW-RL
```

7. Assign the class-of-service rule set to the security policy for logical system LSYS1.

```
user@host# set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy
from_internet match source-address any
user@host# set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy
from_internet match destination-address any
user@host# set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy
from_internet match application any
user@host# set logical-systems LSYS1 security policies from-zone trust to-zone trust policy
p1 match dynamic-application junos:web
user@host# set logical-systems LSYS1 security policies from-zone untrust to-zone trust policy
from_internet then permit application-services application-traffic-control rule-set RS1
```

Results

From configuration mode, confirm your configuration by entering the `show logical-systems LSYS1` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show logical-systems LSYS1
security {
```

```
    policies {
        from-zone untrust to-zone trust {
            policy from_internet {
                match {
                    source-address any;
                    destination-address any;
                    application any;
                }
                then {
                    permit {
                        application-services {
                            application-traffic-control {
                                rule-set RS1;
                            }
                        }
                    }
                }
            }
        }
        from-zone trust to-zone trust {
            policy p1 {
                match {
                    dynamic-application junos:web;
                }
            }
        }
    }
}
class-of-service {
    application-traffic-control {
        rate-limiters HTTP-BW-RL {
            bandwidth-limit 512;
        }
        rule-sets RS1 {
            rule RL1 {
                match {
                    application junos:HTTP;
                }
                then {
                    forwarding-class best-effort;
                    dscp-code-point 001000;
                    loss-priority high;
                    rate-limit {
```

```
                    server-to-client HTTP-BW-RL;
                }
                log;
            }
        }
    }
}
```

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

Verification

IN THIS SECTION

- Verifying the class-of-service application-traffic-control counter | 408
- Verifying the class-of-service application-traffic-control statistics rate-limiter | 409

To confirm that the configuration is working properly, perform the below tasks:

Verifying the class-of-service application-traffic-control counter

Purpose

Verify the class-of-service application-traffic-control counter for logical systems.

Action

To verify the configuration is working properly, enter the `show class-of-service application-traffic-control counter logical-system LSYS1` command.

```
user@host>show class-of-service application-traffic-control counter logical-system LSYS1
Logical System: LSYS1

pic: 0/0
  Counter type                                                Value
 Sessions processed                                           1
 Sessions marked                                              0
```

```
  Sessions honored                                          0
  Sessions rate limited                                     0
  Client-to-server flows rate limited                       0
  Server-to-client flows rate limited                       0
  Session default ruleset hit                               0
  Session ignored no default ruleset                        0
```

Meaning

The output displays AppQoS DSCP marking and honoring statistics based on Layer 7 application classifiers.

Verifying the class-of-service application-traffic-control statistics rate-limiter

Purpose

Verify the class-of-service application-traffic-control statistics rate-limiter for logical systems.

Action

To verify the configuration is working properly, enter the `show class-of-service application-traffic-control statistics rate-limiter logical-system LSYS1` command.

```
user@host>show class-of-service application-traffic-control statistics rate-limiter logical-system LSYS1
Logical System: LSYS1

pic: 0/0
```

Meaning

The output displays AppQoS real-time run information about application rate limiting of current or recent sessions.

Logical Systems in a Chassis Cluster

IN THIS SECTION

- Understanding Logical Systems in the Context of Chassis Cluster | **410**
- Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (Primary Administrators Only) | **411**
- Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (IPv6) (Primary Administrators Only) | **456**

A chassis cluster provides high availability on SRX Series devices where two devices operate as a single device. Chassis cluster includes the synchronization of configuration files and the dynamic runtime session states between the SRX Series devices, which are part of chassis cluster setup. For more information, see the following topics:

Understanding Logical Systems in the Context of Chassis Cluster

The behavior of a *chassis cluster* whose nodes consist of SRX Series devices running logical systems is the same as that of a cluster whose SRX Series nodes in the cluster are not running logical systems. No difference exists between events that cause a node to fail over. In particular, if a link associated with a single logical system fails, then the device fails over to another node in the cluster.

The primary administrator configures the chassis cluster (including both primary and secondary nodes) before he or she creates and configures the logical systems. Each node in the cluster has the same configuration, as is the case for nodes in a cluster not running logical systems. All logical system configurations are synchronized and replicated between both nodes in the cluster.

When you use SRX Series devices running logical systems within a chassis cluster, you must purchase and install the same number of licenses for each node in the chassis cluster. Logical systems licenses pertain to a single chassis, or node, within a chassis cluster and not to the cluster collectively.

Starting with Junos OS Release 12.3X48-D50, when you configure the logical systems within a chassis cluster, if logical systems licenses on backup node are not sufficient when you `commit` the configuration, a warning message is displayed about the number of licenses required on backup node as well, just as on primary node in all the previous releases.

SEE ALSO

Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (Primary Administrators Only) | 411

Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (IPv6) (Primary Administrators Only) | 456

Understanding the Interconnect Logical System and Logical Tunnel Interfaces | 11

Understanding Logical Systems for SRX Series Services Gateways | 5

Chassis Cluster Overview

Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (Primary Administrators Only)

IN THIS SECTION

- Requirements | 411
- Overview | 412
- Configuration | 415
- Verification | 447

This example shows how to configure logical systems in a basic active/passive chassis cluster.

NOTE: The primary administrator configures the chassis cluster and creates logical systems (including an optional interconnect logical system), administrators, and security profiles. Either the primary administrator or the user logical system administrator configures a user logical system. The configuration is synchronized between nodes in the cluster.

Requirements

Before you begin:

- Obtain two SRX Series Services Gateways with identical hardware configurations. See *Example: Configuring an Active/Passive Chassis Cluster on SRX5800 Devices*. This chassis cluster deployment scenario includes the configuration of the SRX Series device for connections to an MX240 edge router and an EX8208 Ethernet Switch.

- Physically connect the two devices (back-to-back for the fabric and control ports) and ensure that they are the same models. You can configure both the fabric and control ports on the SRX5000 line. For the SRX1400 or SRX1500 devices or the SRX3000 line, you can configure the fabric ports only. (Platform support depends on the Junos OS release in your installation.) See *Connecting SRX Series Devices to Create a Chassis Cluster*.
- Set the chassis cluster ID and node ID on each device and reboot the devices to enable clustering. See *Example: Setting the Node ID and Cluster ID for Security Devices in a Chassis Cluster* .

NOTE: For this example, chassis cluster and logical system configuration is performed on the primary (node 0) device at the root level by the primary administrator. Log in to the device as the primary administrator. See "Understanding the Primary Logical Systems and the Primary Administrator Role" on page 23.

NOTE: When you use SRX Series devices running logical systems in a chassis cluster, you must purchase and install the same number of logical system licenses for each node in the chassis cluster. Logical system licenses pertain to a single chassis or node within a chassis cluster and not to the cluster collectively.

Overview

IN THIS SECTION

- Topology | **413**

In this example, the basic active/passive chassis cluster consists of two devices:

- One device actively provides logical systems, along with maintaining control of the chassis cluster.
- The other device passively maintains its state for cluster failover capabilities should the active device become inactive.

NOTE: Logical systems in an active/active chassis cluster are configured in a similar manner as for logical systems in an active/passive chassis cluster. For active/active chassis clusters, there can be multiple redundancy groups that can be primary on different nodes.

The primary administrator configures the following logical systems on the primary device (node 0):

- Primary logical system—The primary administrator configures a security profile to provision portions of the system's security resources to the primary logical system and configures the resources of the primary logical system.
- User logical systems LSYS1 and LSYS2 and their administrators—The primary administrator also configures security profiles to provision portions of the system's security resources to user logical systems. The user logical system administrator can then configure interfaces, routing, and security resources allocated to his or her logical system.
- Interconnect logical system LSYS0 that connects logical systems on the device—The primary administrator configures logical tunnel interfaces between the interconnect logical system and each logical system. These peer interfaces effectively allow for the establishment of tunnels.

NOTE: This example does not describe configuring features such as NAT, IDP, or VPNs for a logical system. See "SRX Series Logical Systems Primary Administrator Configuration Tasks Overview" on page 24 and "User Logical Systems Configuration Overview" on page 51 for more information about features that can be configured for logical systems.

If you are performing proxy ARP in a chassis cluster configuration, you must apply the proxy ARP configuration to the reth interfaces rather than the member interfaces because the reth interfaces contain the logical configurations. See *Configuring Proxy ARP for NAT (CLI Procedure)*.

Topology

Figure 10 on page 414 shows the topology used in this example.

Figure 10: Logical Systems in a Chassis Cluster

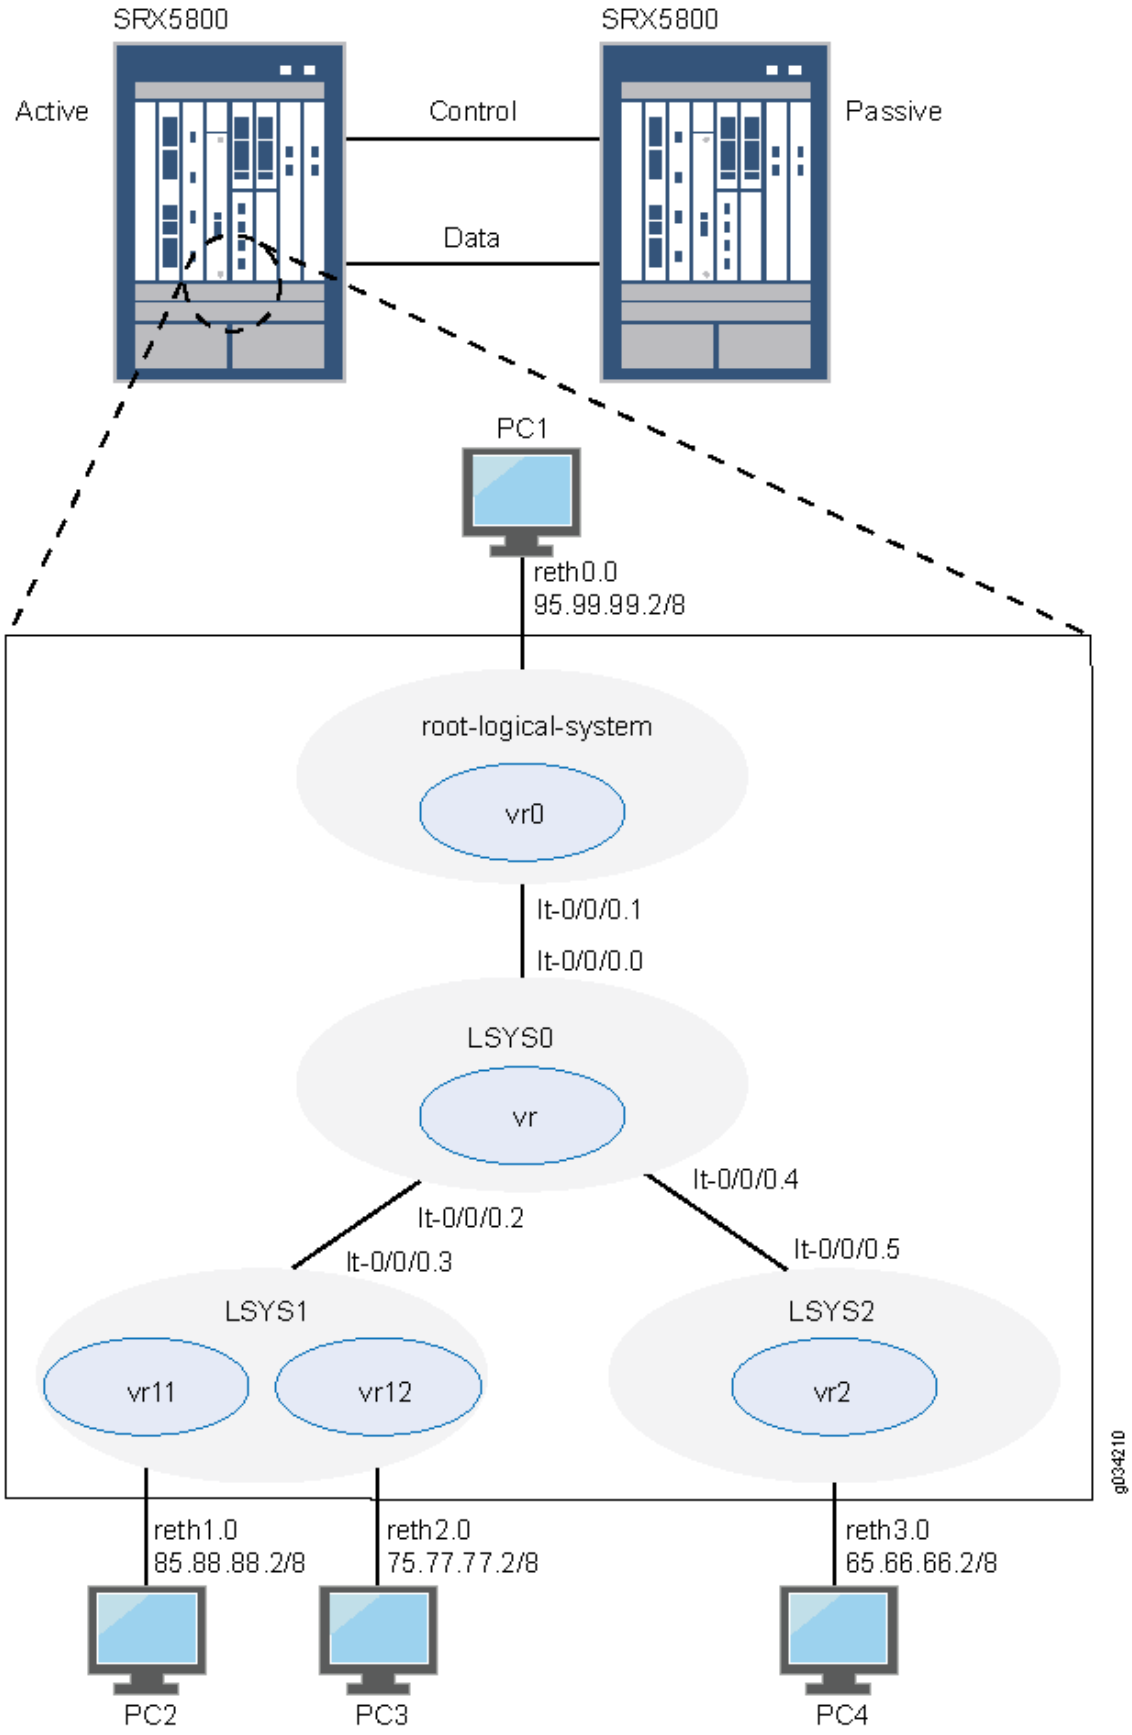

Configuration

IN THIS SECTION

- Chassis Cluster Configuration (Primary Administrator) | 415
- Logical System Configuration (Primary Administrator) | 421
- User Logical System Configuration (User Logical System Administrator) | 434

Chassis Cluster Configuration (Primary Administrator)

CLI Quick Configuration

To quickly create logical systems and user logical system administrators and configure the primary and interconnect logical systems, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the `[edit]` hierarchy level.

On {primary:node0}

```
set chassis cluster control-ports fpc 0 port 0
set chassis cluster control-ports fpc 6 port 0
set interfaces fab0 fabric-options member-interfaces ge-1/1/0
set interfaces fab1 fabric-options member-interfaces ge-7/1/0
set groups node0 system host-name SRX5800-1
set groups node0 interfaces fxp0 unit 0 family inet address 10.157.90.24/9
set groups node0 system backup-router 10.157.64.1 destination 0.0.0.0/0
set groups node1 system host-name SRX5800-2
set groups node1 interfaces fxp0 unit 0 family inet address 10.157.90.23/19
set groups node1 system backup-router 10.157.64.1 destination 0.0.0.0/0
set apply-groups “${node}”
set chassis cluster reth-count 5
set chassis cluster redundancy-group 0 node 0 priority 200
set chassis cluster redundancy-group 0 node 1 priority 100
set chassis cluster redundancy-group 1 node 0 priority 200
set chassis cluster redundancy-group 1 node 1 priority 100
set interfaces ge-1/0/0 gigether-options redundant-parent reth0
set interfaces ge-1/0/1 gigether-options redundant-parent reth1
set interfaces ge-1/0/2 gigether-options redundant-parent reth2
set interfaces ge-1/0/3 gigether-options redundant-parent reth3
```

```
set interfaces ge-7/0/0 gigether-options redundant-parent reth0
set interfaces ge-7/0/1 gigether-options redundant-parent reth1
set interfaces ge-7/0/2 gigether-options redundant-parent reth2
set interfaces ge-7/0/3 gigether-options redundant-parent reth3
set interfaces reth0 redundant-ether-options redundancy-group 1
set interfaces reth0 unit 0 family inet address 95.99.99.1/8
set interfaces reth1 redundant-ether-options redundancy-group 1
set interfaces reth2 redundant-ether-options redundancy-group 1
set interfaces reth3 redundant-ether-options redundancy-group 1
```

Step-by-Step Procedure

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

To configure a chassis cluster:

NOTE: Perform the following steps on the primary device (node 0). They are automatically copied over to the secondary device (node 1) when you execute a `commit` command.

1. Configure control ports for the clusters.

```
[edit chass cluster]
user@host# set control-ports fpc 0 port 0
user@host# set control-ports fpc 6 port 0
```

2. Configure the fabric (data) ports of the cluster that are used to pass RTOs in active/passive mode.

```
[edit interfaces]
user@host# set fab0 fabric-options member-interfaces ge-1/1/0
user@host# set fab1 fabric-options member-interfaces ge-7/1/0
```

3. Assign some elements of the configuration to a specific member. Configure out-of-band management on the fxp0 interface of the SRX Services Gateway using separate IP addresses for the individual control planes of the cluster.

```
[edit]
user@host# set groups node0 system host-name SRX5800-1
user@host# set groups node0 interfaces fxp0 unit 0 family inet address 10.157.90.24/9
user@host# set groups node0 system backup-router 10.157.64.1 destination 0.0.0.0/0
user@host# set groups node1 system host-name SRX5800-2
user@host# set groups node1 interfaces fxp0 unit 0 family inet address 10.157.90.23/19
user@host# set groups node1 system backup-router 10.157.64.1 destination 0.0.0.0/0
user@host# set apply-groups “${node}”
```

4. Configure redundancy groups for chassis clustering.

```
[edit chassis cluster]
user@host# set reth-count 5
user@host# set redundancy-group 0 node 0 priority 200
user@host# set redundancy-group 0 node 1 priority 100
user@host# set redundancy-group 1 node 0 priority 200
user@host# set redundancy-group 1 node 1 priority 100
```

5. Configure the data interfaces on the platform so that in the event of a data plane failover, the other chassis cluster member can take over the connection seamlessly.

```
[edit interfaces]
user@host# set ge-1/0/0 gigether-options redundant-parent reth0
user@host# set ge-1/0/1 gigether-options redundant-parent reth1
user@host# set ge-1/0/2 gigether-options redundant-parent reth2
user@host# set ge-1/0/3 gigether-options redundant-parent reth3
user@host# set ge-7/0/0 gigether-options redundant-parent reth0
user@host# set ge-7/0/1 gigether-options redundant-parent reth1
user@host# set ge-7/0/2 gigether-options redundant-parent reth2
user@host# set ge-7/0/3 gigether-options redundant-parent reth3
user@host# set reth0 redundant-ether-options redundancy-group 1
user@host# set reth0 unit 0 family inet address 95.99.99.1/8
user@host# set reth1 redundant-ether-options redundancy-group 1
user@host# set reth2 redundant-ether-options redundancy-group 1
user@host# set reth3 redundant-ether-options redundancy-group 1
```

Results

From operational mode, confirm your configuration by entering the `show configuration` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host> show configuration
version ;
groups {
    node0 {
        system {
            host-name SRX58001;
            backup-router 10.157.64.1 destination 0.0.0.0/0;

        }
        interfaces {
            fxp0 {
                unit 0 {
                    family inet {
                        address 10.157.90.24/9;
                    }
                }
            }
        }
    }
    node1 {
        system {
            host-name SRX58002;
            backup-router 10.157.64.1 destination 0.0.0.0/0;

        }
        interfaces {
            fxp0 {
                unit 0 {
                    family inet {
                        address 10.157.90.23/19;
                    }
                }
            }
        }
    }
}
```

```
apply-groups "${node}";
chassis {
    cluster {
        control-link-recovery;
        reth-count 5;
        control-ports {
            fpc 0 port 0;
            fpc 6 port 0;
        }
        redundancy-group 0 {
            node 0 priority 200;
            node 1 priority 100;
        }
        redundancy-group 1 {
            node 0 priority 200;
            node 1 priority 100;
        }
    }
}
interfaces {
    ge-1/0/0 {
        gigether-options {
            redundant-parent reth0;
        }
    }
    ge-1/0/1 {
        gigether-options {
            redundant-parent reth1;
        }
    }
    ge-1/0/2 {
        gigether-options {
            redundant-parent reth2;
        }
    }
    ge-1/0/3 {
        gigether-options {
            redundant-parent reth3;
        }
    }
    ge-7/0/0 {
        gigether-options {
            redundant-parent reth0;
```

```
        }
    }
    ge-7/0/1 {
        gigether-options {
            redundant-parent reth1;
        }
    }
    ge-7/0/2 {
        gigether-options {
            redundant-parent reth2;
        }
    }
    ge-7/0/3 {
        gigether-options {
            redundant-parent reth3;
        }
    }
    fab0 {
        fabric-options {
            member-interfaces {
                ge-1/1/0;
            }
        }
    }
    fab1 {
        fabric-options {
            member-interfaces {
                ge-7/1/0;
            }
        }
    }
    reth0 {
        redundant-ether-options {
            redundancy-group 1;
        }
        unit 0 {
            family inet {
                address 95.99.99.1/8;
            }
        }
    }
    reth1 {
        redundant-ether-options {
```

```
            redundancy-group 1;
        }
    }
    reth2 {
        redundant-ether-options {
            redundancy-group 1;
        }
    }
    reth3 {
        redundant-ether-options {
            redundancy-group 1;
        }
    }
}
```

Logical System Configuration (Primary Administrator)

CLI Quick Configuration

To quickly create logical systems and user logical system administrators and configure the primary and interconnect logical systems, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the `[edit]` hierarchy level.

NOTE: You are prompted to enter and then reenter plain-text passwords.

On {primary:node0}

```
set logical-systems LSYS1
set logical-systems LSYS2
set logical-systems LSYS0
set system login class lsys1 logical-system LSYS1
set system login class lsys1 permissions all
set system login user lsys1admin full-name lsys1-admin
set system login user lsys1admin class lsys1
set user lsys1admin authentication plain-text-password
set system login class lsys2 logical-system LSYS2
set system login class lsys2 permissions all
set system login user lsys2admin full-name lsys2-admin
set system login user lsys2admin class lsys2
```

```
set system login user lsys2admin authentication plain-text-password
set system security-profile SP-root policy maximum 200
set system security-profile SP-root policy reserved 100
set system security-profile SP-root zone maximum 200
set system security-profile SP-root zone reserved 100
set system security-profile SP-root flow-session maximum 200
set system security-profile SP-root flow-session reserved 100
set system security-profile SP-root root-logical-system
set system security-profile SP0 logical-system LSYS0
set system security-profile SP1 policy maximum 100
set system security-profile SP1 policy reserved 50
set system security-profile SP1 zone maximum 100
set system security-profile SP1 zone reserved 50
set system security-profile SP1 flow-session maximum 100
set system security-profile SP1 flow-session reserved 50
set system security-profile SP1 logical-system LSYS1
set system security-profile SP2 policy maximum 100
set system security-profile SP2 policy reserved 50
set system security-profile SP2 zone maximum 100
set system security-profile SP2 zone reserved 50
set system security-profile SP2 flow-session maximum 100
set system security-profile SP2 flow-session reserved 50
set system security-profile SP2 logical-system LSYS2
set interfaces lt-0/0/0 unit 1 encapsulation ethernet
set interfaces lt-0/0/0 unit 1 peer-unit 0
set interfaces lt-0/0/0 unit 1 family inet address 2.1.1.1/24
set routing-instances vr0 instance-type virtual-router
set routing-instances vr0 interface lt-0/0/0.1
set routing-instances vr0 interface reth0.0
set routing-instances vr0 routing-options static route 85.0.0.0/8 next-hop 2.1.1.3
set routing-instances vr0 routing-options static route 75.0.0.0/8 next-hop 2.1.1.3
set routing-instances vr0 routing-options static route 65.0.0.0/8 next-hop 2.1.1.5
set security zones security-zone root-trust host-inbound-traffic system-services all
set security zones security-zone root-trust host-inbound-traffic protocols all
set security zones security-zone root-trust interfaces reth0.0
set security zones security-zone root-untrust host-inbound-traffic system-services all
set security zones security-zone root-untrust host-inbound-traffic protocols all
set security zones security-zone root-untrust interfaces lt-0/0/0.1
set security policies from-zone root-trust to-zone root-untrust policy root-Trust_to_root-
Untrust match source-address any
set security policies from-zone root-trust to-zone root-untrust policy root-Trust_to_root-
Untrust match destination-address any
set security policies from-zone root-trust to-zone root-untrust policy root-Trust_to_root-
```

```
Untrust match application any
set security policies from-zone root-trust to-zone root-untrust policy root-Trust_to_root-
Untrust then permit
set security policies from-zone root-untrust to-zone root-trust policy root-Untrust_to_root-
Trust match source-address any
set security policies from-zone root-untrust to-zone root-trust policy root-Untrust_to_root-
Trust match destination-address any
set security policies from-zone root-untrust to-zone root-trust policy root-Untrust_to_root-
Trust match application any
set security policies from-zone root-untrust to-zone root-trust policy root-Untrust_to_root-
Trust then permit
set security policies from-zone root-untrust to-zone root-untrust policy root-Untrust_to_root-
Untrust match source-address any
set security policies from-zone root-untrust to-zone root-untrust policy root-Untrust_to_root-
Untrust match destination-address any
set security policies from-zone root-untrust to-zone root-untrust policy root-Untrust_to_root-
Untrust match application any
set security policies from-zone root-untrust to-zone root-untrust policy root-Untrust_to_root-
Untrust then permit
set security policies from-zone root-trust to-zone root-trust policy root-Trust_to_root-Trust
match source-address any
set security policies from-zone root-trust to-zone root-trust policy root-Trust_to_root-Trust
match destination-address any
set security policies from-zone root-trust to-zone root-trust policy root-Trust_to_root-Trust
match application any
set security policies from-zone root-trust to-zone root-trust policy root-Trust_to_root-Trust
then permit
set logical-systems LSYS0 interfaces lt-0/0/0 unit 0 encapsulation ethernet-vpls
set logical-systems LSYS0 interfaces lt-0/0/0 unit 0 peer-unit 1
set logical-systems LSYS0 interfaces lt-0/0/0 unit 2 encapsulation ethernet-vpls
set logical-systems LSYS0 interfaces lt-0/0/0 unit 2 peer-unit 3
set logical-systems LSYS0 interfaces lt-0/0/0 unit 4 encapsulation ethernet-vpls
set logical-systems LSYS0 interfaces lt-0/0/0 unit 4 peer-unit 5
set logical-systems LSYS0 routing-instances vr instance-type vpls
set logical-systems LSYS0 routing-instances vr interface lt-0/0/0.0
set logical-systems LSYS0 routing-instances vr interface lt-0/0/0.2
set logical-systems LSYS0 routing-instances vr interface lt-0/0/0.4
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 encapsulation ethernet
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 peer-unit 2
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 family inet address 2.1.1.3/24
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 encapsulation ethernet
```

```
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 peer-unit 4
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 family inet address 2.1.1.5/24
```

Step-by-Step Procedure

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

To create logical systems and user logical system administrators and configure the primary and interconnect logical systems:

1. Create the interconnect and user logical systems.

```
[edit logical-systems]
user@host# set LSYS0
user@host# set LSYS1
user@host# set LSYS2
```

2. Configure user logical system administrators.

Step-by-Step Procedure

a. Configure the user logical system administrator for LSYS1.

```
[edit system login]
user@host# set class lsys1 logical-system LSYS1
user@host# set class lsys1 permissions all
user@host# set user lsys1admin full-name lsys1-admin
user@host# set user lsys1admin class lsys1
user@host# set user lsys1admin authentication plain-text-password
```

b. Configure the user logical system administrator for LSYS2.

```
[edit system login]
user@host# set class lsys2 logical-system LSYS2
user@host# set class lsys2 permissions all
user@host# set user lsys2admin full-name lsys2-admin
user@host# set user lsys2admin class lsys2
user@host# set user lsys2admin authentication plain-text-password
```

3. Configure security profiles and assign them to logical systems.

   Step-by-Step Procedure

   a. Configure a security profile and assign it to the root logical system.

```
[edit system security-profile]
user@host# set SP-root policy maximum 200
user@host# set SP-root policy reserved 100
user@host# set SP-root zone maximum 200
user@host# set SP-root zone reserved 100
user@host# set SP-root flow-session maximum 200
user@host# set SP-root flow-session reserved 100
user@host# set SP-root root-logical-system
```

   b. Assign a dummy security profile containing no resources to the interconnect logical system LSYS0.

```
[edit system security-profile]
user@host# set SP0 logical-system LSYS0
```

   c. Configure a security profile and assign it to LSYS1.

```
[edit system security-profile]
user@host# set SP1 policy maximum 100
user@host# set SP1 policy reserved 50
user@host# set SP1 zone maximum 100
user@host# set SP1 zone reserved 50
user@host# set SP1 flow-session maximum 100
user@host# set SP1 flow-session reserved 50
user@host# set SP1 logical-system LSYS1
```

   d. Configure a security profile and assign it to LSYS2.

```
[edit system security-profile]
user@host# set SP2 policy maximum 100
user@host# set SP2 policy reserved 50
user@host# set SP2 zone maximum 100
user@host# set SP2 zone reserved 50
```

```
user@host# set SP2 flow-session maximum 100
user@host# set SP2 flow-session reserved 50
user@host# set SP2 logical-system LSYS2
```

4. Configure the primary logical system.

Step-by-Step Procedure

a. Configure logical tunnel interfaces.

```
[edit interfaces]
user@host# set lt-0/0/0 unit 1 encapsulation ethernet
user@host# set lt-0/0/0 unit 1 peer-unit 0
user@host# set lt-0/0/0 unit 1 family inet address 2.1.1.1/24
```

b. Configure a routing instance.

```
[edit routing-instances]
user@host# set vr0 instance-type virtual-router
user@host# set vr0 interface lt-0/0/0.1
user@host# set vr0 interface reth0.0
user@host# set vr0 routing-options static route 85.0.0.0/8 next-hop 2.1.1.3
user@host# set vr0 routing-options static route 75.0.0.0/8 next-hop 2.1.1.3
user@host# set vr0 routing-options static route 65.0.0.0/8 next-hop 2.1.1.5
```

c. Configure zones.

```
[edit security zones]
user@host# set security-zone root-trust host-inbound-traffic system-services all
user@host# set security-zone root-trust host-inbound-traffic protocols all
user@host# set security-zone root-trust interfaces reth0.0
user@host# set security-zone root-untrust host-inbound-traffic system-services all
user@host# set security-zone root-untrust host-inbound-traffic protocols all
user@host# set security-zone root-untrust interfaces lt-0/0/0.1
```

d. Configure security policies.

```
[edit security policies from-zone root-trust to-zone root-untrust]
user@host# set policy root-Trust_to_root-Untrust match source-address any
user@host# set policy root-Trust_to_root-Untrust match destination-address any
user@host# set policy root-Trust_to_root-Untrust match application any
user@host# set policy root-Trust_to_root-Untrust then permit
```

```
[edit security policies from-zone root-untrust to-zone root-trust]
user@host# set policy root-Untrust_to_root-Trust match source-address any
user@host# set policy root-Untrust_to_root-Trust match destination-address any
user@host# set policy root-Untrust_to_root-Trust match application any
user@host# set policy root-Untrust_to_root-Trust then permit
```

```
[edit security policies from-zone root-untrust to-zone root-untrust]
user@host# set policy root-Untrust_to_root-Untrust match source-address any
user@host# set policy root-Untrust_to_root-Untrust match destination-address any
user@host# set policy root-Untrust_to_root-Untrust match application any
user@host# set policy root-Untrust_to_root-Untrust then permit
```

```
[edit security policies from-zone root-trust to-zone root-trust]
user@host# set policy root-Trust_to_root-Trust match source-address any
user@host# set policy root-Trust_to_root-Trust match destination-address any
user@host# set policy root-Trust_to_root-Trust match application any
user@host# set policy root-Trust_to_root-Trust then permit
```

5. Configure the interconnect logical system.

Step-by-Step Procedure

a. Configure logical tunnel interfaces.

```
[edit logical-systems LSYS0 interfaces]
user@host# set lt-0/0/0 unit 0 encapsulation ethernet-vpls
user@host# set lt-0/0/0 unit 0 peer-unit 1
user@host# set lt-0/0/0 unit 2 encapsulation ethernet-vpls
```

```
user@host# set lt-0/0/0 unit 2 peer-unit 3
user@host# set lt-0/0/0 unit 4 encapsulation ethernet-vpls
user@host# set lt-0/0/0 unit 4 peer-unit 5
```

b. Configure the VPLS routing instance.

```
[edit logical-systems LSYS0 routing-instances]
user@host# set vr instance-type vpls
user@host# set vr interface lt-0/0/0.0
user@host# set vr interface lt-0/0/0.2
user@host# set vr interface lt-0/0/0.4
```

6. Configure logical tunnel interfaces for the user logical systems.

Step-by-Step Procedure

a. Configure logical tunnel interfaces for LSYS1.

```
[edit logical-systems LSYS1 interfaces ]
user@host# set lt-0/0/0 unit 3 encapsulation ethernet
user@host# set lt-0/0/0 unit 3 peer-unit 2
user@host# set lt-0/0/0 unit 3 family inet address 2.1.1.3/24
```

b. Configure logical tunnel interfaces for LSYS2.

```
[edit logical-systems LSYS2 interfaces ]
user@host# set lt-0/0/0 unit 5 encapsulation ethernet
user@host# set lt-0/0/0 unit 5 peer-unit 4
user@host# set lt-0/0/0 unit 5 family inet address 2.1.1.5/24
```

Results

From configuration mode, confirm the configuration for LSYS0 by entering the `show logical-systems LSYS0` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show logical-systems LSYS0
interfaces {
    lt-0/0/0 {
        unit 0 {
            encapsulation ethernet-vpls;
            peer-unit 1;
        }
        unit 2 {
            encapsulation ethernet-vpls;
            peer-unit 3;
        }
        unit 4 {
            encapsulation ethernet-vpls;
            peer-unit 5;
        }
    }
}
    routing-instances {
        vr {
            instance-type vpls;
            interface lt-0/0/0.0;
            interface lt-0/0/0.2;
            interface lt-0/0/0.4;
        }
    }
```

From configuration mode, confirm the configuration for the primary logical system by entering the `show interfaces`, `show routing-instances`, and `show security` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
lt-0/0/0 {
    unit 1 {
```

```
        encapsulation ethernet;
        peer-unit 0;
        family inet {
            address 2.1.1.1/24;
        }
    }
}
    ge-1/0/0 {
        gigether-options {
            redundant-parent reth0;
        }
    }
    ge-1/0/1 {
        gigether-options {
            redundant-parent reth1;
        }
    }
    ge-1/0/2 {
        gigether-options {
            redundant-parent reth2;
        }
    }
    ge-1/0/3 {
        gigether-options {
            redundant-parent reth3;
        }
    }
    ge-7/0/0 {
        gigether-options {
            redundant-parent reth0;
        }
    }
    ge-7/0/1 {
        gigether-options {
            redundant-parent reth1;
        }
    }
    ge-7/0/2 {
        gigether-options {
            redundant-parent reth2;
        }
    }
    ge-7/0/3 {
```

```
        gigether-options {
            redundant-parent reth3;
        }
    }
    fab0 {
        fabric-options {
            member-interfaces {
                ge-1/1/0;
            }
        }
    }
    fab1 {
        fabric-options {
            member-interfaces {
                ge-7/1/0;
            }
        }
    }
    reth0 {
        redundant-ether-options {
            redundancy-group 1;
        }
        unit 0 {
            family inet {
                address 95.99.99.1/8;
            }
        }
    }
    reth1 {
        redundant-ether-options {
            redundancy-group 1;
        }
    }
    reth2 {
        redundant-ether-options {
            redundancy-group 1;
        }
    }
    reth3 {
        redundant-ether-options {
            redundancy-group 1;
        }
    }
```

```
[edit]
user@host# show routing-instances
vr0 {
    instance-type virtual-router;
    interface lt-0/0/0.1;
    interface reth0.0;
    routing-options {
        static {
            route 85.0.0.0/8 next-hop 2.1.1.3;
            route 75.0.0.0/8 next-hop 2.1.1.3;
            route 65.0.0.0/8 next-hop 2.1.1.5;
        }
    }
}
[edit]
user@host# show security
policies {
    from-zone root-trust to-zone root-untrust {
        policy root-Trust_to_root-Untrust {
            match {
                source-address any;
                destination-address any;
                application any;
            }
            then {
                permit;
            }
        }
    }
    from-zone root-untrust to-zone root-trust {
        policy root-Untrust_to_root-Trust {
            match {
                source-address any;
                destination-address any;
                application any;
            }
            then {
                permit;
            }
        }
    }
    from-zone root-untrust to-zone root-untrust {
        policy root-Untrust_to_root-Untrust {
```

```
            match {
                source-address any;
                destination-address any;
                application any;
            }
            then {
                permit;
            }
        }
    }
    from-zone root-trust to-zone root-trust {
        policy root-Trust_to_root-Trust {
            match {
                source-address any;
                destination-address any;
                application any;
            }
            then {
                permit;
            }
        }
    }
}
    zones {
        security-zone root-trust {
            host-inbound-traffic {
                system-services {
                    all;
                }
                protocols {
                    all;
                }
            }
            interfaces {
                reth0.0;
            }
        }
        security-zone root-untrust {
            host-inbound-traffic {
                system-services {
                    all;
                }
                protocols {
```

```
                    all;
                }
            }
            interfaces {
                lt-0/0/0.1;
            }
        }
    }
```

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

User Logical System Configuration (User Logical System Administrator)

CLI Quick Configuration

To quickly configure user logical systems, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the `[edit]` hierarchy level.

Enter the following commands while logged in as the user logical system administrator for LSYS1:

```
set interfaces reth1 unit 0 family inet address 85.88.88.1/8
set interfaces reth2 unit 0 family inet address 75.77.77.1/8
set routing-instances vr11 instance-type virtual-router
set routing-instances vr11 interface lt-0/0/0.3
set routing-instances vr11 interface reth1.0
set routing-instances vr11 routing-options static route 65.0.0.0/8 next-hop 2.1.1.5
set routing-instances vr11 routing-options static route 95.0.0.0/8 next-hop 2.1.1.1
set routing-instances vr12 instance-type virtual-router
set routing-instances vr12 interface reth2.0
set routing-instances vr12 routing-options interface-routes rib-group inet vr11vr12v4
set routing-instances vr12 routing-options static route 85.0.0.0/8 next-table vr11.inet.0
set routing-instances vr12 routing-options static route 95.0.0.0/8 next-table vr11.inet.0
set routing-instances vr12 routing-options static route 65.0.0.0/8 next-table vr11.inet.0
set routing-instances vr12 routing-options static route 2.1.1.0/24 next-table vr11.inet.0
set routing-options rib-groups vr11vr12v4 import-rib vr11.inet.0
set routing-options rib-groups vr11vr12v4 import-rib vr12.inet.0
set security zones security-zone lsys1-trust host-inbound-traffic system-services all
set security zones security-zone lsys1-trust host-inbound-traffic protocols all
set security zones security-zone lsys1-trust interfaces reth1.0
set security zones security-zone lsys1-trust interfaces lt-0/0/0.3
set security zones security-zone lsys1-untrust host-inbound-traffic system-services all
```

```
set security zones security-zone lsys1-untrust host-inbound-traffic protocols all
set security zones security-zone lsys1-untrust interfaces reth2.0
set security policies from-zone lsys1-trust to-zone lsys1-untrust policy lsys1trust-to-
lsys1untrust match source-address any
set security policies from-zone lsys1-trust to-zone lsys1-untrust policy lsys1trust-to-
lsys1untrust match destination-address any
set security policies from-zone lsys1-trust to-zone lsys1-untrust policy lsys1trust-to-
lsys1untrust match application any
set security policies from-zone lsys1-trust to-zone lsys1-untrust policy lsys1trust-to-
lsys1untrust then permit
set security policies from-zone lsys1-untrust to-zone lsys1-trust policy lsys1untrust-to-
lsys1trust match source-address any
set security policies from-zone lsys1-untrust to-zone lsys1-trust policy lsys1untrust-to-
lsys1trust match destination-address any
set security policies from-zone lsys1-untrust to-zone lsys1-trust policy lsys1untrust-to-
lsys1trust match application any
set security policies from-zone lsys1-untrust to-zone lsys1-trust policy lsys1untrust-to-
lsys1trust then permit
set security policies from-zone lsys1-untrust to-zone lsys1-untrust policy lsys1untrust-to-
lsys1untrust match source-address any
set security policies from-zone lsys1-untrust to-zone lsys1-untrust policy lsys1untrust-to-
lsys1untrust match destination-address any
set security policies from-zone lsys1-untrust to-zone lsys1-untrust policy lsys1untrust-to-
lsys1untrust match application any
set security policies from-zone lsys1-untrust to-zone lsys1-untrust policy lsys1untrust-to-
lsys1untrust then permit
set security policies from-zone lsys1-trust to-zone lsys1-trust policy lsys1trust-to-lsys1trust
match source-address any
set security policies from-zone lsys1-trust to-zone lsys1-trust policy lsys1trust-to-lsys1trust
match destination-address any
set security policies from-zone lsys1-trust to-zone lsys1-trust policy lsys1trust-to-lsys1trust
match application any
set security policies from-zone lsys1-trust to-zone lsys1-trust policy lsys1trust-to-lsys1trust
then permit
```

Enter the following commands while logged in as the user logical system administrator for LSYS2:

```
set interfaces reth3 unit 0 family inet address 65.66.66.1/8
set routing-instances vr2 instance-type virtual-router
set routing-instances vr2 interface lt-0/0/0.5
set routing-instances vr2 interface reth3.0
set routing-instances vr2 routing-options static route 75.0.0.0/8 next-hop 2.1.1.3
```

```
set routing-instances vr2 routing-options static route 85.0.0.0/8 next-hop 2.1.1.3
set routing-instances vr2 routing-options static route 95.0.0.0/8 next-hop 2.1.1.1
set security zones security-zone lsys2-trust host-inbound-traffic system-services all
set security zones security-zone lsys2-trust host-inbound-traffic protocols all
set security zones security-zone lsys2-trust interfaces reth3.0
set security zones security-zone lsys2-untrust host-inbound-traffic system-services all
set security zones security-zone lsys2-untrust host-inbound-traffic protocols all
set security zones security-zone lsys2-untrust interfaces lt-0/0/0.5
set security policies from-zone lsys2-trust to-zone lsys2-untrust policy lsys2trust-to-
lsys2untrust match source-address any
set security policies from-zone lsys2-trust to-zone lsys2-untrust policy lsys2trust-to-
lsys2untrust match destination-address any
set security policies from-zone lsys2-trust to-zone lsys2-untrust policy lsys2trust-to-
lsys2untrust match application any
set security policies from-zone lsys2-trust to-zone lsys2-untrust policy lsys2trust-to-
lsys2untrust then permit
set security policies from-zone lsys2-untrust to-zone lsys2-trust policy lsys2untrust-to-
lsys2trust match source-address any
set security policies from-zone lsys2-untrust to-zone lsys2-trust policy lsys2untrust-to-
lsys2trust match destination-address any
set security policies from-zone lsys2-untrust to-zone lsys2-trust policy lsys2untrust-to-
lsys2trust match application any
set security policies from-zone lsys2-untrust to-zone lsys2-trust policy lsys2untrust-to-
lsys2trust then permit
set security policies from-zone lsys2-untrust to-zone lsys2-untrust policy lsys2untrust-to-
lsys2untrust match source-address any
set security policies from-zone lsys2-untrust to-zone lsys2-untrust policy lsys2untrust-to-
lsys2untrust match destination-address any
set security policies from-zone lsys2-untrust to-zone lsys2-untrust policy lsys2untrust-to-
lsys2untrust match application any
set security policies from-zone lsys2-untrust to-zone lsys2-untrust policy lsys2untrust-to-
lsys2untrust then permit
set security policies from-zone lsys2-trust to-zone lsys2-trust policy lsys2trust-to-lsys2trust
match source-address any
set security policies from-zone lsys2-trust to-zone lsys2-trust policy lsys2trust-to-lsys2trust
match destination-address any
set security policies from-zone lsys2-trust to-zone lsys2-trust policy lsys2trust-to-lsys2trust
match application any
set security policies from-zone lsys2-trust to-zone lsys2-trust policy lsys2trust-to-lsys2trust
then permit
```

Step-by-Step Procedure

NOTE: The user logical system administrator performs the following configuration while logged in to his or her user logical system. The primary administrator can also configure a user logical system at the [`edit logical-systems` *`logical-system`*] hierarchy level.

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

To configure the LSYS1 user logical system:

1. Configure interfaces.

```
[edit interfaces]
lsys1-admin@host:LSYS1# set reth1 unit 0 family inet address 85.88.88.1/8
lsys1-admin@host:LSYS1# set reth2 unit 0 family inet address 75.77.77.1/8
```

2. Configure routing.

```
[edit routing-instances]
lsys1-admin@host:LSYS1# set vr11 instance-type virtual-router
lsys1-admin@host:LSYS1# set vr11 interface lt-0/0/0.3
lsys1-admin@host:LSYS1# set vr11 interface reth1.0
lsys1-admin@host:LSYS1# set vr11 routing-options static route 65.0.0.0/8 next-hop 2.1.1.5
lsys1-admin@host:LSYS1# set vr11 routing-options static route 95.0.0.0/8 next-hop 2.1.1.1
lsys1-admin@host:LSYS1# set vr12 instance-type virtual-router
lsys1-admin@host:LSYS1# set vr12 interface reth2.0
lsys1-admin@host:LSYS1# set vr12 routing-options interface-routes rib-group inet vr11vr12v4
lsys1-admin@host:LSYS1# set vr12 routing-options static route 85.0.0.0/8 next-table
vr11.inet.0
lsys1-admin@host:LSYS1# set vr12 routing-options static route 95.0.0.0/8 next-table
vr11.inet.0
lsys1-admin@host:LSYS1# set vr12 routing-options static route 65.0.0.0/8 next-table
vr11.inet.0
```

```
lsys1-admin@host:LSYS1# set vr12 routing-options static route 2.1.1.0/24 next-table
vr11.inet.0
```

```
[edit routing-options]
lsys1-admin@host:LSYS1# set rib-groups vr11vr12v4 import-rib vr11.inet.0
lsys1-admin@host:LSYS1# set rib-groups vr11vr12v4 import-rib vr12.inet.0
```

3. Configure zones and security policies.

```
[edit security zones]
lsys1-admin@host:LSYS1# set security-zone lsys1-trust host-inbound-traffic system-services all
lsys1-admin@host:LSYS1# set security-zone lsys1-trust host-inbound-traffic protocols all
lsys1-admin@host:LSYS1# set security-zone lsys1-trust interfaces reth1.0
lsys1-admin@host:LSYS1# set security-zone lsys1-trust interfaces lt-0/0/0.3
lsys1-admin@host:LSYS1# set security-zone lsys1-untrust host-inbound-traffic system-services
all
lsys1-admin@host:LSYS1# set security-zone lsys1-untrust host-inbound-traffic protocols all
lsys1-admin@host:LSYS1# set security-zone lsys1-untrust interfaces reth2.0
```

```
[edit security policies from-zone lsys1-trust to-zone lsys1-untrust]
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1untrust match source-address any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1untrust match destination-address any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1untrust match application any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1untrust then permit
```

```
[edit security policies from-zone lsys1-untrust to-zone lsys1-trust]
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1trust match source-address any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1trust match destination-address any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1trust match application any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1trust then permit
```

```
[edit security policies from-zone lsys1-untrust to-zone lsys1-untrust]
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1untrust match source-address any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1untrust match destination-address any
```

```
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1untrust match application any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1untrust then permit
```

```
[edit security policies from-zone lsys1-trust to-zone lsys1-trust]
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1trust match source-address any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1trust match destination-address any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1trust match application any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1trust then permit
```

Step-by-Step Procedure

To configure the LSYS2 user logical system:

1. Configure interfaces.

```
[edit interfaces]
lsys2-admin@host:LSYS2# set reth3 unit 0 family inet address 65.66.66.1/8
```

2. Configure routing.

```
[edit routing-instances]
lsys2-admin@host:LSYS2# set vr2 instance-type virtual-router
lsys2-admin@host:LSYS2# set vr2 interface lt-0/0/0.5
lsys2-admin@host:LSYS2# set vr2 interface reth3.0
lsys2-admin@host:LSYS2# set vr2 routing-options static route 75.0.0.0/8 next-hop 2.1.1.3
lsys2-admin@host:LSYS2# set vr2 routing-options static route 85.0.0.0/8 next-hop 2.1.1.3
lsys2-admin@host:LSYS2# set vr2 routing-options static route 95.0.0.0/8 next-hop 2.1.1.1
```

3. Configure zones and security policies.

```
[edit security zones]
lsys2-admin@host:LSYS2# set security-zone lsys2-trust host-inbound-traffic system-services all
lsys2-admin@host:LSYS2# set security-zone lsys2-trust host-inbound-traffic protocols all
lsys2-admin@host:LSYS2# set security-zone lsys2-trust interfaces reth3.0
lsys2-admin@host:LSYS2# set security zones security-zone lsys2-untrust host-inbound-traffic
system-services all
```

```
lsys2-admin@host:LSYS2# set security-zone lsys2-untrust host-inbound-traffic protocols all
lsys2-admin@host:LSYS2# set security-zone lsys2-untrust interfaces lt-0/0/0.5
```

```
[edit security policies from-zone lsys2-trust to-zone lsys2-untrust]
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2untrust match source-address any
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2untrust match destination-address any
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2untrust match application any
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2untrust then permit
```

```
[edit security policies from-zone from-zone lsys2-untrust to-zone lsys2-trust]
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2trust match source-address any
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2trust match destination-address any
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2trust match application any
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2trust then permit
```

```
[edit security policies from-zone lsys2-untrust to-zone lsys2-untrust]
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2untrust match source-address any
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2untrust match destination-address any
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2untrust match application any
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2untrust then permit
```

```
[edit security policies from-zone lsys2-trust to-zone lsys2-trust]
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2trust match source-address any
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2trust match destination-address any
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2trust match application any
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2trust then permit
```

Results

From configuration mode, confirm the configuration for LSYS1 by entering the `show interfaces`, `show routing-instances`, `show routing-options`, and `show security` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
lsys1-admin@host:LSYS1# show interfaces
interfaces {
```

```
    lt-0/0/0 {
        unit 3 {
            encapsulation ethernet;
            peer-unit 2;
            family inet {
                address 2.1.1.3/24;
            }
        }
    }
    reth1 {
        unit 0 {
            family inet {
                address 85.88.88.1/8;
            }
        }
    }
    reth2 {
        unit 0 {
            family inet {
                address 75.77.77.1/8;
            }
        }
    }
}
[edit]
lsys1-admin@host:LSYS1# show routing-instances
routing-instances {
    vr11 {
        instance-type virtual-router;
        interface lt-0/0/0.3;
        interface reth1.0;
        routing-options {
            static {
                route 65.0.0.0/8 next-hop 2.1.1.5;
                route 95.0.0.0/8 next-hop 2.1.1.1;
            }
        }
    }
    vr12 {
        instance-type virtual-router;
        interface reth2.0;
        routing-options {
            interface-routes {
```

```
                rib-group inet vr11vr12v4;
            }
            static {
                route 85.0.0.0/8 next-table vr11.inet.0;
                route 95.0.0.0/8 next-table vr11.inet.0;
                route 65.0.0.0/8 next-table vr11.inet.0;
                route 2.1.1.0/24 next-table vr11.inet.0;
            }
        }
    }
}
[edit]
lsys1-admin@host:LSYS1# show routing-options
rib-groups {
    vr11vr12v4 {
        import-rib [ vr11.inet.0 vr12.inet.0 ];
    }
}
[edit]
lsys1-admin@host:LSYS1# show security
security {
    policies {
        from-zone lsys1-trust to-zone lsys1-untrust {
            policy lsys1trust-to-lsys1untrust {
                match {
                    source-address any;
                    destination-address any;
                    application any;
                }
                then {
                    permit;
                }
            }
        }
        from-zone lsys1-untrust to-zone lsys1-trust {
            policy lsys1untrust-to-lsys1trust {
                match {
                    source-address any;
                    destination-address any;
                    application any;
                }
                then {
                    permit;
```

```
                }
            }
        }
        from-zone lsys1-untrust to-zone lsys1-untrust {
            policy lsys1untrust-to-lsys1untrust {
                match {
                    source-address any;
                    destination-address any;
                    application any;
                }
                then {
                    permit;
                }
            }
        }
        from-zone lsys1-trust to-zone lsys1-trust {
            policy lsys1trust-to-lsys1trust {
                match {
                    source-address any;
                    destination-address any;
                    application any;
                }
                then {
                    permit;
                }
            }
        }
    }
    zones {
        security-zone lsys1-trust {
            host-inbound-traffic {
                system-services {
                    all;
                }
                protocols {
                    all;
                }
            }
            interfaces {
                reth1.0;
                lt-0/0/0.3;
            }
        }
```

```
        security-zone lsys1-untrust {
            host-inbound-traffic {
                system-services {
                    all;
                }
                protocols {
                    all;
                }
            }
            interfaces {
                reth2.0;
            }
        }
    }
}
```

From configuration mode, confirm the configuration for LSYS2 by entering the `show interfaces`, `show routing-instances`, and `show security` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
lsys2-admin@host:LSYS2# show interfaces
[edit]
    interfaces {
    lt-0/0/0 {
        unit 5 {
            encapsulation ethernet;
            peer-unit 4;
            family inet {
                address 2.1.1.5/24;
            }
        }
    }
    reth3 {
        unit 0 {
            family inet {
                address 65.66.66.1/8;
            }
        }
    }
}
[edit]
lsys2-admin@host:LSYS2# show routing-instances
```

```
routing-instances {
    vr2 {
        instance-type virtual-router;
        interface lt-0/0/0.5;
        interface reth3.0;
        routing-options {
            static {
                route 75.0.0.0/8 next-hop 2.1.1.3;
                route 85.0.0.0/8 next-hop 2.1.1.3;
                route 95.0.0.0/8 next-hop 2.1.1.1;
            }
        }
    }
}
[edit]
lsys2-admin@host:LSYS2# show security
security {
    policies {
        from-zone lsys2-trust to-zone lsys2-untrust {
            policy lsys2trust-to-lsys2untrust {
                match {
                    source-address any;
                    destination-address any;
                    application any;
                }
                then {
                    permit;
                }
            }
        }
        from-zone lsys2-untrust to-zone lsys2-trust {
            policy lsys2untrust-to-lsys2trust {
                match {
                    source-address any;
                    destination-address any;
                    application any;
                }
                then {
                    permit;
                }
            }
        }
        from-zone lsys2-untrust to-zone lsys2-untrust {
```

```
            policy lsys2untrust-to-lsys2untrust {
                match {
                    source-address any;
                    destination-address any;
                    application any;
                }
                then {
                    permit;
                }
            }
        }
        from-zone lsys2-trust to-zone lsys2-trust {
            policy lsys2trust-to-lsys2trust {
                match {
                    source-address any;
                    destination-address any;
                    application any;
                }
                then {
                    permit;
                }
            }
        }
    }
    zones {
        security-zone lsys2-trust {
            host-inbound-traffic {
                system-services {
                    all;
                }
                protocols {
                    all;
                }
            }
            interfaces {
                reth3.0;
            }
        }
        security-zone lsys2-untrust {
            host-inbound-traffic {
                system-services {
                    all;
                }
```

```
                protocols {
                    all;
                }
            }
            interfaces {
                lt-0/0/0.5;
            }
        }
    }
}
```

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

Verification

IN THIS SECTION

- Verifying Chassis Cluster Status | **447**
- Troubleshooting Chassis Cluster with Logs | **448**
- Verifying Logical System Licenses | **448**
- Verifying Logical System License Usage | **449**
- Verifying Intra-Logical System Traffic on a Logical System | **450**
- Verifying Intra-Logical System Traffic Within All Logical Systems | **451**
- Verifying Traffic Between User Logical Systems | **452**

Confirm that the configuration is working properly.

Verifying Chassis Cluster Status

Purpose

Verify the chassis cluster status, failover status, and redundancy group information.

Action

From operational mode, enter the `show chassis cluster status` command.

```
{primary:node0}
show chassis cluster status
Cluster ID: 1
Node                  Priority          Status    Preempt  Manual failover

Redundancy group: 0 , Failover count: 1
    node0                   200         primary        no       no
    node1                   100         secondary      no       no

Redundancy group: 1 , Failover count: 1
    node0                   200         primary        no       no
    node1                   100         secondary      no       no
```

Troubleshooting Chassis Cluster with Logs

Purpose

Identify any chassis cluster issues by looking at the logs on both nodes.

Action

From operational mode, enter these `show log` commands.

```
user@host> show log jsrpd
user@host> show log chassisd
user@host> show log messages
user@host> show log dcd
user@host> show traceoptions
```

Verifying Logical System Licenses

Purpose

Verify information about logical system licenses.

Action

From operational mode, enter the `show system license status logical-system all` command.

```
{primary:node0}
user@host> show system license status logical-system all
node0:
--------------------------------------------------------------------------
Logical system license status:

logical system name                license status
root-logical-system                enabled
LSYS0                              enabled
LSYS1                              enabled
LSYS2                              enabled
```

Verifying Logical System License Usage

Purpose

Verify information about logical system license usage.

NOTE: The actual number of licenses used is only displayed on the primary node.

Action

From operational mode, enter the `show system license` command.

```
{primary:node0}
user@host> show system license
License usage:
                                Licenses     Licenses    Licenses    Expiry
  Feature name                       used    installed      needed
  logical-system                        4           25           0    permanent

Licenses installed:
  License identifier: JUNOS305013
  License version: 2
```

```
  Valid for device: JN110B54BAGB
  Features:
    logical-system-25 - Logical System Capacity
      permanent
```

Verifying Intra-Logical System Traffic on a Logical System

Purpose

Verify information about currently active security sessions within a logical system.

Action

From operational mode, enter the `show security flow session logical-system LSYS1` command.

```
{primary:node0}
user@host> show security flow session logical-system LSYS1
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:
Total sessions: 0

Flow Sessions on FPC2 PIC1:

Session ID: 90000114, Policy name: lsys1trust-to-lsys1untrust/8, State: Active, Timeout: 1782,
Valid
  In: 85.88.88.2/34538 --> 75.77.77.2/23;tcp, If: reth1.0, Pkts: 33, Bytes: 1881
  Out: 75.77.77.2/23 --> 85.88.88.2/34538;tcp, If: reth2.0, Pkts: 28, Bytes: 2329
Total sessions: 1

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

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:
```

```
Total sessions: 0

Flow Sessions on FPC2 PIC1:

Session ID: 90000001, Policy name: lsys1trust-to-lsys1untrust/8, State: Backup, Timeout: 14388,
Valid
  In: 85.88.88.2/34538 --> 75.77.77.2/23;tcp, If: reth1.0, Pkts: 0, Bytes: 0
  Out: 75.77.77.2/23 --> 85.88.88.2/34538;tcp, If: reth2.0, Pkts: 0, Bytes: 0
Total sessions: 1
```

Verifying Intra-Logical System Traffic Within All Logical Systems

Purpose

Verify information about currently active security sessions on all logical systems.

Action

From operational mode, enter the `show security flow session logical-system all` command.

```
{primary:node0}
user@host> show security flow session logical-system all
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:
Total sessions: 0

Flow Sessions on FPC2 PIC1:

Session ID: 90000114, Policy name: lsys1trust-to-lsys1untrust/8, State: Active, Timeout: 1776,
Valid
Logical system: LSYS1
  In: 85.88.88.2/34538 --> 75.77.77.2/23;tcp, If: reth1.0, Pkts: 33, Bytes: 1881
  Out: 75.77.77.2/23 --> 85.88.88.2/34538;tcp, If: reth2.0, Pkts: 28, Bytes: 2329
Total sessions: 1

node1:
```

```
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:
Total sessions: 0

Flow Sessions on FPC2 PIC1:

Session ID: 90000001, Policy name: lsys1trust-to-lsys1untrust/8, State: Backup, Timeout: 14382,
Valid
Logical system: LSYS1
  In: 85.88.88.2/34538 --> 75.77.77.2/23;tcp, If: reth1.0, Pkts: 0, Bytes: 0
  Out: 75.77.77.2/23 --> 85.88.88.2/34538;tcp, If: reth2.0, Pkts: 0, Bytes: 0
Total sessions: 1
```

Verifying Traffic Between User Logical Systems

Purpose

Verify information about currently active security sessions between logical systems.

Action

From operational mode, enter the `show security flow session logical-system` *`logical-system-name`* command.

```
{primary:node0}
user@host> show security flow session logical-system LSYS1

node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000094, Policy name: root-Untrust_to_root-Trust/5, State: Active, Timeout: 1768,
Valid
  In: 75.77.77.2/34590 --> 95.99.99.2/23;tcp, If: lt-0/0/0.1, Pkts: 23, Bytes: 1351
  Out: 95.99.99.2/23 --> 75.77.77.2/34590;tcp, If: reth0.0, Pkts: 22, Bytes: 1880
Total sessions: 1
```

```
Flow Sessions on FPC2 PIC0:
Total sessions: 0

Flow Sessions on FPC2 PIC1:
Total sessions: 0

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

Flow Sessions on FPC0 PIC1:

Session ID: 10000002, Policy name: root-Untrust_to_root-Trust/5, State: Backup, Timeout: 14384,
Valid
  In: 75.77.77.2/34590 --> 95.99.99.2/23;tcp, If: lt-0/0/0.1, Pkts: 0, Bytes: 0
  Out: 95.99.99.2/23 --> 75.77.77.2/34590;tcp, If: reth0.0, Pkts: 0, Bytes: 0
Total sessions: 1

Flow Sessions on FPC2 PIC0:
Total sessions: 0

Flow Sessions on FPC2 PIC1:
Total sessions: 0
```

```
{primary:node0}
user@host> show security flow session logical-system LSYS2

node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:

Session ID: 80000089, Policy name: lsys2untrust-to-lsys2trust/13, State: Active, Timeout: 1790,
Valid
  In: 85.88.88.2/34539 --> 65.66.66.2/23;tcp, If: lt-0/0/0.5, Pkts: 40, Bytes: 2252
  Out: 65.66.66.2/23 --> 85.88.88.2/34539;tcp, If: reth3.0, Pkts: 32, Bytes: 2114
Total sessions: 1

Flow Sessions on FPC2 PIC1:
```

```
Total sessions: 0

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

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:

Session ID: 80000002, Policy name: lsys2untrust-to-lsys2trust/13, State: Backup, Timeout: 14398,
Valid
  In: 85.88.88.2/34539 --> 65.66.66.2/23;tcp, If: lt-0/0/0.5, Pkts: 0, Bytes: 0
  Out: 65.66.66.2/23 --> 85.88.88.2/34539;tcp, If: reth3.0, Pkts: 0, Bytes: 0
Total sessions: 1

Flow Sessions on FPC2 PIC1:
Total sessions: 0
```

```
{primary:node0}
user@host> show security flow session logical-system all

node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:

Session ID: 80000088, Policy name: lsys1trust-to-lsys1trust/11, State: Active, Timeout: 1782,
Valid
Logical system: LSYS1
  In: 85.88.88.2/34539 --> 65.66.66.2/23;tcp, If: reth1.0, Pkts: 40, Bytes: 2252
  Out: 65.66.66.2/23 --> 85.88.88.2/34539;tcp, If: lt-0/0/0.3, Pkts: 32, Bytes: 2114

Session ID: 80000089, Policy name: lsys2untrust-to-lsys2trust/13, State: Active, Timeout: 1782,
Valid
Logical system: LSYS2
  In: 85.88.88.2/34539 --> 65.66.66.2/23;tcp, If: lt-0/0/0.5, Pkts: 40, Bytes: 2252
  Out: 65.66.66.2/23 --> 85.88.88.2/34539;tcp, If: reth3.0, Pkts: 32, Bytes: 2114
```

```
Total sessions: 2

Flow Sessions on FPC2 PIC1:
Total sessions: 0

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

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:

Session ID: 80000001, Policy name: lsys1trust-to-lsys1trust/11, State: Backup, Timeout: 14382,
Valid
Logical system: LSYS1
  In: 85.88.88.2/34539 --> 65.66.66.2/23;tcp, If: reth1.0, Pkts: 0, Bytes: 0
  Out: 65.66.66.2/23 --> 85.88.88.2/34539;tcp, If: lt-0/0/0.3, Pkts: 0, Bytes: 0

Session ID: 80000002, Policy name: lsys2untrust-to-lsys2trust/13, State: Backup, Timeout: 14390,
Valid
Logical system: LSYS2
  In: 85.88.88.2/34539 --> 65.66.66.2/23;tcp, If: lt-0/0/0.5, Pkts: 0, Bytes: 0
  Out: 65.66.66.2/23 --> 85.88.88.2/34539;tcp, If: reth3.0, Pkts: 0, Bytes: 0
Total sessions: 2

Flow Sessions on FPC2 PIC1:
Total sessions: 0
```

SEE ALSO

- Understanding Logical Systems in the Context of Chassis Cluster | **410**
- Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (IPv6) (Primary Administrators Only) | **456**
- Example: Configuring an Active/Passive Chassis Cluster on SRX5800 Devices
- Chassis Cluster Overview

Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (IPv6) (Primary Administrators Only)

IN THIS SECTION

- Requirements | 456
- Overview | 457
- Configuration | 460
- Verification | 492

This example shows how to configure logical systems in a basic active/passive chassis cluster with IPv6 addresses.

NOTE: The primary administrator configures the chassis cluster and creates logical systems (including an optional interconnect logical system), administrators, and security profiles. Either the primary administrator or the user logical system administrator configures a user logical system. The configuration is synchronized between nodes in the cluster.

Requirements

Before you begin:

- Obtain two SRX Series Services Gateways with identical hardware configurations. See *Example: Configuring an Active/Passive Chassis Cluster on SRX5800 Devices*. This chassis cluster deployment scenario includes the configuration of the SRX Series device for connections to an MX240 edge router and an EX8208 Ethernet Switch.
- Physically connect the two devices (back-to-back for the fabric and control ports) and ensure that they are the same models. You can configure both the fabric and control ports on the SRX5000 line. For the SRX1400 or SRX1500 devices or the SRX3000 line, you can configure the fabric ports only. (Platform support depends on the Junos OS release in your installation.)
- Set the chassis cluster ID and node ID on each device and reboot the devices to enable clustering. See *Example: Setting the Node ID and Cluster ID for Security Devices in a Chassis Cluster* .

NOTE: For this example, chassis cluster and logical system configuration is performed on the primary (node 0) device at the root level by the primary administrator. Log in to the device as the primary administrator. See "Understanding the Primary Logical Systems and the Primary Administrator Role" on page 23.

NOTE: When you use SRX Series devices running logical systems in a chassis cluster, you must purchase and install the same number of logical system licenses for each node in the chassis cluster. Logical system licenses pertain to a single chassis or node within a chassis cluster and not to the cluster collectively.

Overview

IN THIS SECTION

- Topology | **458**

In this example, the basic active/passive chassis cluster consists of two devices:

- One device actively provides logical systems, along with maintaining control of the chassis cluster.
- The other device passively maintains its state for cluster failover capabilities should the active device become inactive.

NOTE: Logical systems in an active/active chassis cluster are configured in a similar manner as for logical systems in an active/passive chassis cluster. For active/active chassis clusters, there can be multiple redundancy groups that can be primary on different nodes.

The primary administrator configures the following logical systems on the primary device (node 0):

- Primary logical system—The primary administrator configures a security profile to provision portions of the system's security resources to the primary logical system and configures the resources of the primary logical system.
- User logical systems LSYS1 and LSYS2 and their administrators—The primary administrator also configures security profiles to provision portions of the system's security resources to user logical

systems. The user logical system administrator can then configure interfaces, routing, and security resources allocated to his or her logical system.

- Interconnect logical system LSYS0 that connects logical systems on the device—The primary administrator configures logical tunnel interfaces between the interconnect logical system and each logical system. These peer interfaces effectively allow for the establishment of tunnels.

> **NOTE**: This example does not describe configuring features such as NAT, IDP, or VPNs for a logical system. See "SRX Series Logical Systems Primary Administrator Configuration Tasks Overview" on page 24 and "User Logical Systems Configuration Overview" on page 51 for more information about features that can be configured for logical systems.
>
> If you are performing proxy ARP in a chassis cluster configuration, you must apply the proxy ARP configuration to the reth interfaces rather than the member interfaces because the reth interfaces contain the logical configurations. See *Configuring Proxy ARP for NAT (CLI Procedure)*.

Topology

Figure 11 on page 459 shows the topology used in this example.

Figure 11: Logical Systems in a Chassis Cluster (IPv6)

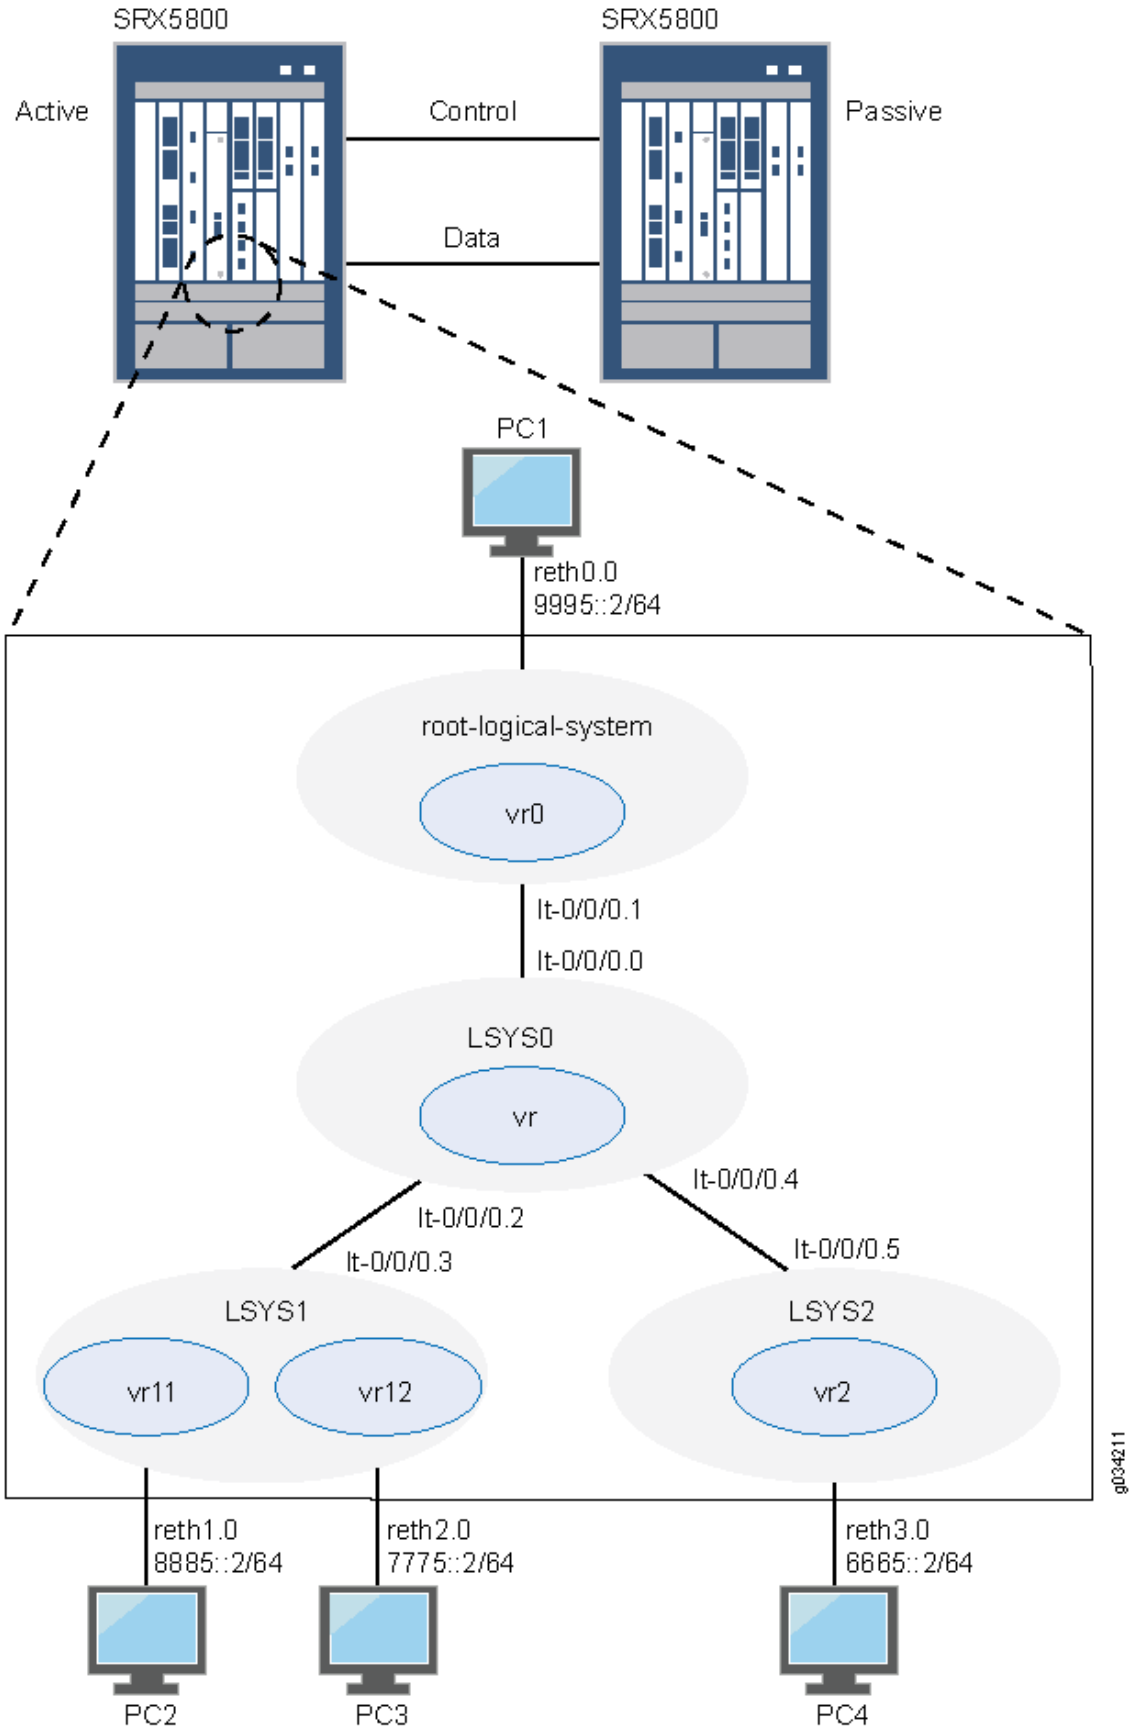

Configuration

IN THIS SECTION

- Chassis Cluster Configuration with IPv6 Addresses (Primary Administrator) | 460
- Logical System Configuration with IPv6 Addresses (Primary Administrator) | 466
- User Logical System Configuration with IPv6 (User Logical System Administrator) | 479

Chassis Cluster Configuration with IPv6 Addresses (Primary Administrator)

CLI Quick Configuration

To quickly create logical systems and user logical system administrators and configure the primary and interconnect logical systems, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the `[edit]` hierarchy level.

On {primary:node0}

```
set chassis cluster control-ports fpc 0 port 0
set chassis cluster control-ports fpc 6 port 0
set interfaces fab0 fabric-options member-interfaces ge-1/1/0
set interfaces fab1 fabric-options member-interfaces ge-7/1/0
set groups node0 system host-name SRX5800-1
set groups node0 interfaces fxp0 unit 0 family inet address 10.157.90.24/9
set groups node0 system backup-router 10.157.64.1 destination 0.0.0.0/0
set groups node1 system host-name SRX5800-2
set groups node1 interfaces fxp0 unit 0 family inet address 10.157.90.23/19
set groups node1 system backup-router 10.157.64.1 destination 0.0.0.0/0
set apply-groups “${node}”
set chassis cluster reth-count 5
set chassis cluster redundancy-group 0 node 0 priority 200
set chassis cluster redundancy-group 0 node 1 priority 100
set chassis cluster redundancy-group 1 node 0 priority 200
set chassis cluster redundancy-group 1 node 1 priority 100
set interfaces ge-1/0/0 gigether-options redundant-parent reth0
set interfaces ge-1/0/1 gigether-options redundant-parent reth1
set interfaces ge-1/0/2 gigether-options redundant-parent reth2
set interfaces ge-1/0/3 gigether-options redundant-parent reth3
```

```
set interfaces ge-7/0/0 gigether-options redundant-parent reth0
set interfaces ge-7/0/1 gigether-options redundant-parent reth1
set interfaces ge-7/0/2 gigether-options redundant-parent reth2
set interfaces ge-7/0/3 gigether-options redundant-parent reth3
set interfaces reth0 redundant-ether-options redundancy-group 1
set interfaces reth0 unit 0 family inet6 address 9995::1/64
set interfaces reth1 redundant-ether-options redundancy-group 1
set interfaces reth2 redundant-ether-options redundancy-group 1
set interfaces reth3 redundant-ether-options redundancy-group 1
```

Step-by-Step Procedure

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

To configure a chassis cluster:

NOTE: Perform the following steps on the primary device (node 0). They are automatically copied over to the secondary device (node 1) when you execute a `commit` command.

1. Configure control ports for the clusters.

```
[edit chassis cluster]
user@host# set control-ports fpc 0 port 0
user@host# set control-ports fpc 6 port 0
```

2. Configure the fabric (data) ports of the cluster that are used to pass RTOs in active/passive mode.

```
[edit interfaces]
user@host# set fab0 fabric-options member-interfaces ge-1/1/0
user@host# set fab1 fabric-options member-interfaces ge-7/1/0
```

3. Assign some elements of the configuration to a specific member. Configure out-of-band management on the fxp0 interface of the SRX Services Gateway using separate IP addresses for the individual control planes of the cluster.

```
[edit]
user@host# set groups node0 system host-name SRX5800-1
user@host# set groups node0 interfaces fxp0 unit 0 family inet address 10.157.90.24/9
user@host# set groups node0 system backup-router 10.157.64.1 destination 0.0.0.0/0
user@host# set groups node1 system host-name SRX5800-2
user@host# set groups node1 interfaces fxp0 unit 0 family inet address 10.157.90.23/19
user@host# set groups node1 system backup-router 10.157.64.1 destination 0.0.0.0/0
user@host# set apply-groups “${node}”
```

4. Configure redundancy groups for chassis clustering.

```
[edit chassis cluster]
user@host# set reth-count 5
user@host# set redundancy-group 0 node 0 priority 200
user@host# set redundancy-group 0 node 1 priority 100
user@host# set redundancy-group 1 node 0 priority 200
user@host# set redundancy-group 1 node 1 priority 100
```

5. Configure the data interfaces on the platform so that in the event of a data plane failover, the other chassis cluster member can take over the connection seamlessly.

```
[edit interfaces]
user@host# set ge-1/0/0 gigether-options redundant-parent reth0
user@host# set ge-1/0/1 gigether-options redundant-parent reth1
user@host# set ge-1/0/2 gigether-options redundant-parent reth2
user@host# set ge-1/0/3 gigether-options redundant-parent reth3
user@host# set ge-7/0/0 gigether-options redundant-parent reth0
user@host# set ge-7/0/1 gigether-options redundant-parent reth1
user@host# set ge-7/0/2 gigether-options redundant-parent reth2
user@host# set ge-7/0/3 gigether-options redundant-parent reth3
user@host# set reth0 redundant-ether-options redundancy-group 1
user@host# set reth0 unit 0 family inet6 address 9995::1/64
user@host# set reth1 redundant-ether-options redundancy-group 1
user@host# set reth2 redundant-ether-options redundancy-group 1
user@host# set reth3 redundant-ether-options redundancy-group 1
```

Results

From operational mode, confirm your configuration by entering the `show configuration` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host> show configuration
version ;
groups {
    node0 {
        system {
            host-name SRX58001;
            backup-router 10.157.64.1 destination 0.0.0.0/0;

        }
        interfaces {
            fxp0 {
                unit 0 {
                    family inet {
                        address 10.157.90.24/9;
                    }
                }
            }
        }
    }
    node1 {
        system {
            host-name SRX58002;
            backup-router 10.157.64.1 destination 0.0.0.0/0;

        }
        interfaces {
            fxp0 {
                unit 0 {
                    family inet {
                        address 10.157.90.23/19;
                    }
                }
            }
        }
    }
}
```

```
apply-groups "${node}";
chassis {
    cluster {
        control-link-recovery;
        reth-count 5;
        control-ports {
            fpc 0 port 0;
            fpc 6 port 0;
        }
        redundancy-group 0 {
            node 0 priority 200;
            node 1 priority 100;
        }
        redundancy-group 1 {
            node 0 priority 200;
            node 1 priority 100;
        }
    }
}
interfaces {
    ge-1/0/0 {
        gigether-options {
            redundant-parent reth0;
        }
    }
    ge-1/0/1 {
        gigether-options {
            redundant-parent reth1;
        }
    }
    ge-1/0/2 {
        gigether-options {
            redundant-parent reth2;
        }
    }
    ge-1/0/3 {
        gigether-options {
            redundant-parent reth3;
        }
    }
    ge-7/0/0 {
        gigether-options {
            redundant-parent reth0;
```

```
        }
    }
    ge-7/0/1 {
        gigether-options {
            redundant-parent reth1;
        }
    }
    ge-7/0/2 {
        gigether-options {
            redundant-parent reth2;
        }
    }
    ge-7/0/3 {
        gigether-options {
            redundant-parent reth3;
        }
    }
    fab0 {
        fabric-options {
            member-interfaces {
                ge-1/1/0;
            }
        }
    }
    fab1 {
        fabric-options {
            member-interfaces {
                ge-7/1/0;
            }
        }
    }
    reth0 {
        redundant-ether-options {
            redundancy-group 1;
        }
        unit 0 {
            family inet6 {
                address 9995::1/64;
            }
        }
    }
    reth1 {
        redundant-ether-options {
```

```
            redundancy-group 1;
        }
    }
    reth2 {
        redundant-ether-options {
            redundancy-group 1;
        }
    }
    reth3 {
        redundant-ether-options {
            redundancy-group 1;
        }
    }
}
```

Logical System Configuration with IPv6 Addresses (Primary Administrator)

CLI Quick Configuration

To quickly create logical systems and user logical system administrators and configure the primary and interconnect logical systems, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the `[edit]` hierarchy level.

NOTE: You are prompted to enter and then reenter plain-text passwords.

On {primary:node0}

```
set logical-systems LSYS1
set logical-systems LSYS2
set logical-systems LSYS0
set system login class lsys1 logical-system LSYS1
set system login class lsys1 permissions all
set system login user lsys1admin full-name lsys1-admin
set system login user lsys1admin class lsys1
set user lsys1admin authentication plain-text-password
set system login class lsys2 logical-system LSYS2
set system login class lsys2 permissions all
set system login user lsys2admin full-name lsys2-admin
set system login user lsys2admin class lsys2
```

```
set system login user lsys2admin authentication plain-text-password
set system security-profile SP-root policy maximum 200
set system security-profile SP-root policy reserved 100
set system security-profile SP-root zone maximum 200
set system security-profile SP-root zone reserved 100
set system security-profile SP-root flow-session maximum 200
set system security-profile SP-root flow-session reserved 100
set system security-profile SP-root root-logical-system
set system security-profile SP0 logical-system LSYS0
set system security-profile SP1 policy maximum 100
set system security-profile SP1 policy reserved 50
set system security-profile SP1 zone maximum 100
set system security-profile SP1 zone reserved 50
set system security-profile SP1 flow-session maximum 100
set system security-profile SP1 flow-session reserved 50
set system security-profile SP1 logical-system LSYS1
set system security-profile SP2 policy maximum 100
set system security-profile SP2 policy reserved 50
set system security-profile SP2 zone maximum 100
set system security-profile SP2 zone reserved 50
set system security-profile SP2 flow-session maximum 100
set system security-profile SP2 flow-session reserved 50
set system security-profile SP2 logical-system LSYS2
set interfaces lt-0/0/0 unit 1 encapsulation ethernet
set interfaces lt-0/0/0 unit 1 peer-unit 0
set interfaces lt-0/0/0 unit 1 family inet6 address 2111::1/64
set routing-instances vr0 instance-type virtual-router
set routing-instances vr0 interface lt-0/0/0.1
set routing-instances vr0 interface reth0.0
set routing-instances vr0 routing-options rib vr0.inet6.0 static route 8885::/64 next-hop 2111::3
set routing-instances vr0 routing-options rib vr0.inet6.0 static route 7775::/64 next-hop 2111::3
set routing-instances vr0 routing-options rib vr0.inet6.0 static route 6665::/64 next-hop
2111::5
set security zones security-zone root-trust host-inbound-traffic system-services all
set security zones security-zone root-trust host-inbound-traffic protocols all
set security zones security-zone root-trust interfaces reth0.0
set security zones security-zone root-untrust host-inbound-traffic system-services all
set security zones security-zone root-untrust host-inbound-traffic protocols all
set security zones security-zone root-untrust interfaces lt-0/0/0.1
set security policies from-zone root-trust to-zone root-untrust policy root-Trust_to_root-
Untrust match source-address any
set security policies from-zone root-trust to-zone root-untrust policy root-Trust_to_root-
Untrust match destination-address any
```

```
set security policies from-zone root-trust to-zone root-untrust policy root-Trust_to_root-
Untrust match application any
set security policies from-zone root-trust to-zone root-untrust policy root-Trust_to_root-
Untrust then permit
set security policies from-zone root-untrust to-zone root-trust policy root-Untrust_to_root-
Trust match source-address any
set security policies from-zone root-untrust to-zone root-trust policy root-Untrust_to_root-
Trust match destination-address any
set security policies from-zone root-untrust to-zone root-trust policy root-Untrust_to_root-
Trust match application any
set security policies from-zone root-untrust to-zone root-trust policy root-Untrust_to_root-
Trust then permit
set security policies from-zone root-untrust to-zone root-untrust policy root-Untrust_to_root-
Untrust match source-address any
set security policies from-zone root-untrust to-zone root-untrust policy root-Untrust_to_root-
Untrust match destination-address any
set security policies from-zone root-untrust to-zone root-untrust policy root-Untrust_to_root-
Untrust match application any
set security policies from-zone root-untrust to-zone root-untrust policy root-Untrust_to_root-
Untrust then permit
set security policies from-zone root-trust to-zone root-trust policy root-Trust_to_root-Trust
match source-address any
set security policies from-zone root-trust to-zone root-trust policy root-Trust_to_root-Trust
match destination-address any
set security policies from-zone root-trust to-zone root-trust policy root-Trust_to_root-Trust
match application any
set security policies from-zone root-trust to-zone root-trust policy root-Trust_to_root-Trust
then permit
set logical-systems LSYS0 interfaces lt-0/0/0 unit 0 encapsulation ethernet-vpls
set logical-systems LSYS0 interfaces lt-0/0/0 unit 0 peer-unit 1
set logical-systems LSYS0 interfaces lt-0/0/0 unit 2 encapsulation ethernet-vpls
set logical-systems LSYS0 interfaces lt-0/0/0 unit 2 peer-unit 3
set logical-systems LSYS0 interfaces lt-0/0/0 unit 4 encapsulation ethernet-vpls
set logical-systems LSYS0 interfaces lt-0/0/0 unit 4 peer-unit 5
set logical-systems LSYS0 routing-instances vr instance-type vpls
set logical-systems LSYS0 routing-instances vr interface lt-0/0/0.0
set logical-systems LSYS0 routing-instances vr interface lt-0/0/0.2
set logical-systems LSYS0 routing-instances vr interface lt-0/0/0.4
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 encapsulation ethernet
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 peer-unit 2
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 family inet6 address 2111::3/64
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 encapsulation ethernet
```

```
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 peer-unit 4
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 family inet6 address 2111::5/64
```

Step-by-Step Procedure

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

To create logical systems and user logical system administrators and configure the primary and interconnect logical systems:

1. Create the interconnect and user logical systems.

```
[edit logical-systems]
user@host# set LSYS0
user@host# set LSYS1
user@host# set LSYS2
```

2. Configure user logical system administrators.

Step-by-Step Procedure

a. Configure the user logical system administrator for LSYS1.

```
[edit system login]
user@host# set class lsys1 logical-system LSYS1
user@host# set class lsys1 permissions all
user@host# set user lsys1admin full-name lsys1-admin
user@host# set user lsys1admin class lsys1
user@host# set user lsys1admin authentication plain-text-password
```

b. Configure the user logical system administrator for LSYS2.

```
[edit system login]
user@host# set class lsys2 logical-system LSYS2
user@host# set class lsys2 permissions all
user@host# set user lsys2admin full-name lsys2-admin
user@host# set user lsys2admin class lsys2
user@host# set user lsys2admin authentication plain-text-password
```

3. Configure security profiles and assign them to logical systems.

   Step-by-Step Procedure

   a. Configure a security profile and assign it to the root logical system.

```
[edit system security-profile]
user@host# set SP-root policy maximum 200
user@host# set SP-root policy reserved 100
user@host# set SP-root zone maximum 200
user@host# set SP-root zone reserved 100
user@host# set SP-root flow-session maximum 200
user@host# set SP-root flow-session reserved 100
user@host# set SP-root root-logical-system
```

   b. Assign a dummy security profile containing no resources to the interconnect logical system LSYS0.

```
[edit system security-profile]
user@host# set SP0 logical-system LSYS0
```

   c. Configure a security profile and assign it to LSYS1.

```
[edit system security-profile]
user@host# set SP1 policy maximum 100
user@host# set SP1 policy reserved 50
user@host# set SP1 zone maximum 100
user@host# set SP1 zone reserved 50
user@host# set SP1 flow-session maximum 100
user@host# set SP1 flow-session reserved 50
user@host# set SP1 logical-system LSYS1
```

   d. Configure a security profile and assign it to LSYS2.

```
[edit system security-profile]
user@host# set SP2 policy maximum 100
user@host# set SP2 policy reserved 50
user@host# set SP2 zone maximum 100
user@host# set SP2 zone reserved 50
```

```
user@host# set SP2 flow-session maximum 100
user@host# set SP2 flow-session reserved 50
user@host# set SP2 logical-system LSYS2
```

4. Configure the primary logical system.

Step-by-Step Procedure

a. Configure logical tunnel interfaces.

```
[edit interfaces]
user@host# set lt-0/0/0 unit 1 encapsulation ethernet
user@host# set lt-0/0/0 unit 1 peer-unit 0
user@host# set lt-0/0/0 unit 1 family inet6 address 2111::1/64
```

b. Configure a routing instance.

```
[edit routing-instances]
user@host# set vr0 instance-type virtual-router
user@host# set vr0 interface lt-0/0/0.1
user@host# set vr0 interface reth0.0
user@host# set vr0 routing-options rib vr0.inet6.0 static route 8885::/64 next-hop 2111::3
user@host# set vr0 routing-options rib vr0.inet6.0 static route 7775::/64 next-hop 2111::3
user@host# set vr0 routing-options rib vr0.inet6.0 static route 6665::/64 next-hop 2111::5
```

c. Configure zones.

```
[edit security zones]
user@host# set security-zone root-trust host-inbound-traffic system-services all
user@host# set security-zone root-trust host-inbound-traffic protocols all
user@host# set security-zone root-trust interfaces reth0.0
user@host# set security-zone root-untrust host-inbound-traffic system-services all
user@host# set security-zone root-untrust host-inbound-traffic protocols all
user@host# set security-zone root-untrust interfaces lt-0/0/0.1
```

d. Configure security policies.

```
[edit security policies from-zone root-trust to-zone root-untrust]
user@host# set policy root-Trust_to_root-Untrust match source-address any
user@host# set policy root-Trust_to_root-Untrust match destination-address any
user@host# set policy root-Trust_to_root-Untrust match application any
user@host# set policy root-Trust_to_root-Untrust then permit
```

```
[edit security policies from-zone root-untrust to-zone root-trust]
user@host# set policy root-Untrust_to_root-Trust match source-address any
user@host# set policy root-Untrust_to_root-Trust match destination-address any
user@host# set policy root-Untrust_to_root-Trust match application any
user@host# set policy root-Untrust_to_root-Trust then permit
```

```
[edit security policies from-zone root-untrust to-zone root-untrust]
user@host# set policy root-Untrust_to_root-Untrust match source-address any
user@host# set policy root-Untrust_to_root-Untrust match destination-address any
user@host# set policy root-Untrust_to_root-Untrust match application any
user@host# set policy root-Untrust_to_root-Untrust then permit
```

```
[edit security policies from-zone root-trust to-zone root-trust]
user@host# set policy root-Trust_to_root-Trust match source-address any
user@host# set policy root-Trust_to_root-Trust match destination-address any
user@host# set policy root-Trust_to_root-Trust match application any
user@host# set policy root-Trust_to_root-Trust then permit
```

5. Configure the interconnect logical system.

Step-by-Step Procedure

a. Configure logical tunnel interfaces.

```
[edit logical-systems LSYS0 interfaces]
user@host# set lt-0/0/0 unit 0 encapsulation ethernet-vpls
user@host# set lt-0/0/0 unit 0 peer-unit 1
user@host# set lt-0/0/0 unit 2 encapsulation ethernet-vpls
```

```
user@host# set lt-0/0/0 unit 2 peer-unit 3
user@host# set lt-0/0/0 unit 4 encapsulation ethernet-vpls
user@host# set lt-0/0/0 unit 4 peer-unit 5
```

b. Configure the VPLS routing instance.

```
[edit logical-systems LSYS0 routing-instances]
user@host# set vr instance-type vpls
user@host# set vr interface lt-0/0/0.0
user@host# set vr interface lt-0/0/0.2
user@host# set vr interface lt-0/0/0.4
```

6. Configure logical tunnel interfaces for the user logical systems.

Step-by-Step Procedure

a. Configure logical tunnel interfaces for LSYS1.

```
[edit logical-systems LSYS1 interfaces ]
user@host# set lt-0/0/0 unit 3 encapsulation ethernet
user@host# set lt-0/0/0 unit 3 peer-unit 2
user@host# set lt-0/0/0 unit 3 family inet6 address 2111::3/64
```

b. Configure logical tunnel interfaces for LSYS2.

```
[edit logical-systems LSYS2 interfaces ]
user@host# set lt-0/0/0 unit 5 encapsulation ethernet
user@host# set lt-0/0/0 unit 5 peer-unit 4
user@host# set lt-0/0/0 unit 5 family inet6 address 2111::5/64
```

Results

From configuration mode, confirm the configuration for LSYS0 by entering the `show logical-systems LSYS0` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show logical-systems LSYS0
interfaces {
    lt-0/0/0 {
        unit 0 {
            encapsulation ethernet-vpls;
            peer-unit 1;
        }
        unit 2 {
            encapsulation ethernet-vpls;
            peer-unit 3;
        }
        unit 4 {
            encapsulation ethernet-vpls;
            peer-unit 5;
        }
    }
}
    routing-instances {
        vr {
            instance-type vpls;
            interface lt-0/0/0.0;
            interface lt-0/0/0.2;
            interface lt-0/0/0.4;
        }
    }
```

From configuration mode, confirm the configuration for the primary logical system by entering the `show interfaces`, `show routing-instances`, and `show security` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
lt-0/0/0 {
    unit 1 {
```

```
        encapsulation ethernet;
        peer-unit 0;
        family inet6 {
            address 2111::1/64;
        }
    }
}
    ge-1/0/0 {
        gigether-options {
            redundant-parent reth0;
        }
    }
    ge-1/0/1 {
        gigether-options {
            redundant-parent reth1;
        }
    }
    ge-1/0/2 {
        gigether-options {
            redundant-parent reth2;
        }
    }
    ge-1/0/3 {
        gigether-options {
            redundant-parent reth3;
        }
    }
    ge-7/0/0 {
        gigether-options {
            redundant-parent reth0;
        }
    }
    ge-7/0/1 {
        gigether-options {
            redundant-parent reth1;
        }
    }
    ge-7/0/2 {
        gigether-options {
            redundant-parent reth2;
        }
    }
    ge-7/0/3 {
```

```
        gigether-options {
            redundant-parent reth3;
        }
    }
    fab0 {
        fabric-options {
            member-interfaces {
                ge-1/1/0;
            }
        }
    }
    fab1 {
        fabric-options {
            member-interfaces {
                ge-7/1/0;
            }
        }
    }
    reth0 {
        redundant-ether-options {
            redundancy-group 1;
        }
        unit 0 {
            family inet6 {
                address 9995::1/64;
            }
        }
    }
    reth1 {
        redundant-ether-options {
            redundancy-group 1;
        }
    }
    reth2 {
        redundant-ether-options {
            redundancy-group 1;
        }
    }
    reth3 {
        redundant-ether-options {
            redundancy-group 1;
        }
    }
```

```
[edit]
user@host# show routing-instances
vr0 {
    instance-type virtual-router;
    interface lt-0/0/0.1;
    interface reth0.0;
    routing-options {
        rib vr0.inet6.0 {
            static {
                route 8885::/64 next-hop 2111::3;
                route 7775::/64 next-hop 2111::3;
                route 6665::/64 next-hop 2111::5;
            }
        }
    }
}
[edit]
user@host# show security
policies {
    from-zone root-trust to-zone root-untrust {
        policy root-Trust_to_root-Untrust {
            match {
                source-address any;
                destination-address any;
                application any;
            }
            then {
                permit;
            }
        }
    }
    from-zone root-untrust to-zone root-trust {
        policy root-Untrust_to_root-Trust {
            match {
                source-address any;
                destination-address any;
                application any;
            }
            then {
                permit;
            }
        }
    }
```

```
    from-zone root-untrust to-zone root-untrust {
        policy root-Untrust_to_root-Untrust {
            match {
                source-address any;
                destination-address any;
                application any;
            }
            then {
                permit;
            }
        }
    }
    from-zone root-trust to-zone root-trust {
        policy root-Trust_to_root-Trust {
            match {
                source-address any;
                destination-address any;
                application any;
            }
            then {
                permit;
            }
        }
    }
}
    zones {
        security-zone root-trust {
            host-inbound-traffic {
                system-services {
                    all;
                }
                protocols {
                    all;
                }
            }
            interfaces {
                reth0.0;
            }
        }
        security-zone root-untrust {
            host-inbound-traffic {
                system-services {
                    all;
```

```
                }
                protocols {
                    all;
                }
            }
            interfaces {
                lt-0/0/0.1;
            }
        }
    }
```

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

User Logical System Configuration with IPv6 (User Logical System Administrator)

CLI Quick Configuration

To quickly configure user logical systems, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the `[edit]` hierarchy level.

Enter the following commands while logged in as the user logical system administrator for LSYS1:

```
set interfaces reth1 unit 0 family inet6 address 8885::1/64
set interfaces reth2 unit 0 family inet6 address 7775::1/64
set routing-instances vr11 instance-type virtual-router
set routing-instances vr11 interface lt-0/0/0.3
set routing-instances vr11 interface reth1.0
set routing-instances vr11 routing-options rib vr11.inet6.0 static route 6665::/64 next-hop
2111::5
set routing-instances vr11 routing-options rib vr11.inet6.0 static route 9995::/64 next-hop
2111::1
set routing-instances vr12 instance-type virtual-router
set routing-instances vr12 interface reth2.0
set routing-instances vr12 routing-options interface-routes rib-group inet6 vr11vr12v6
set routing-instances vr12 rrouting-options rib vr12.inet6.0 static route 8885::/64 next-table
vr11.inet6.0
set routing-instances vr12 routing-options rib vr12.inet6.0 static route 9995::/64 next-table
vr11.inet6.0
set routing-instances vr12 routing-options rib vr12.inet6.0 static route 6665::/64 next-table
vr11.inet6.0
set routing-instances vr12 routing-options rib vr12.inet6.0 static route 2111::/64 next-table
```

```
vr11.inet6.0
set routing-options rib-groups vr11vr12v6 import-rib vr11.inet6.0
set routing-options rib-groups vr11vr12v6 import-rib vr12.inet6.0
set security zones security-zone lsys1-trust host-inbound-traffic system-services all
set security zones security-zone lsys1-trust host-inbound-traffic protocols all
set security zones security-zone lsys1-trust interfaces reth1.0
set security zones security-zone lsys1-trust interfaces lt-0/0/0.3
set security zones security-zone lsys1-untrust host-inbound-traffic system-services all
set security zones security-zone lsys1-untrust host-inbound-traffic protocols all
set security zones security-zone lsys1-untrust interfaces reth2.0
set security policies from-zone lsys1-trust to-zone lsys1-untrust policy lsys1trust-to-
lsys1untrust match source-address any
set security policies from-zone lsys1-trust to-zone lsys1-untrust policy lsys1trust-to-
lsys1untrust match destination-address any
set security policies from-zone lsys1-trust to-zone lsys1-untrust policy lsys1trust-to-
lsys1untrust match application any
set security policies from-zone lsys1-trust to-zone lsys1-untrust policy lsys1trust-to-
lsys1untrust then permit
set security policies from-zone lsys1-untrust to-zone lsys1-trust policy lsys1untrust-to-
lsys1trust match source-address any
set security policies from-zone lsys1-untrust to-zone lsys1-trust policy lsys1untrust-to-
lsys1trust match destination-address any
set security policies from-zone lsys1-untrust to-zone lsys1-trust policy lsys1untrust-to-
lsys1trust match application any
set security policies from-zone lsys1-untrust to-zone lsys1-trust policy lsys1untrust-to-
lsys1trust then permit
set security policies from-zone lsys1-untrust to-zone lsys1-untrust policy lsys1untrust-to-
lsys1untrust match source-address any
set security policies from-zone lsys1-untrust to-zone lsys1-untrust policy lsys1untrust-to-
lsys1untrust match destination-address any
set security policies from-zone lsys1-untrust to-zone lsys1-untrust policy lsys1untrust-to-
lsys1untrust match application any
set security policies from-zone lsys1-untrust to-zone lsys1-untrust policy lsys1untrust-to-
lsys1untrust then permit
set security policies from-zone lsys1-trust to-zone lsys1-trust policy lsys1trust-to-lsys1trust
match source-address any
set security policies from-zone lsys1-trust to-zone lsys1-trust policy lsys1trust-to-lsys1trust
match destination-address any
set security policies from-zone lsys1-trust to-zone lsys1-trust policy lsys1trust-to-lsys1trust
match application any
set security policies from-zone lsys1-trust to-zone lsys1-trust policy lsys1trust-to-lsys1trust
then permit
```

Enter the following commands while logged in as the user logical system administrator for LSYS2:

```
set interfaces reth3 unit 0 family inet6 address 6665::1/64
set routing-instances vr2 instance-type virtual-router
set routing-instances vr2 interface lt-0/0/0.5
set routing-instances vr2 interface reth3.0
set routing-instances vr2 routing-options rib vr2.inet6.0 static route 7775::/64 next-hop 2111::3
set routing-instances vr2 routing-options rib vr2.inet6.0 static route 8885::/64 next-hop 2111::3
set routing-instances vr2 routing-options rib vr2.inet6.0 static route 9995::/64 next-hop 2111::1
set security zones security-zone lsys2-trust host-inbound-traffic system-services all
set security zones security-zone lsys2-trust host-inbound-traffic protocols all
set security zones security-zone lsys2-trust interfaces reth3.0
set security zones security-zone lsys2-untrust host-inbound-traffic system-services all
set security zones security-zone lsys2-untrust host-inbound-traffic protocols all
set security zones security-zone lsys2-untrust interfaces lt-0/0/0.5
set security policies from-zone lsys2-trust to-zone lsys2-untrust policy lsys2trust-to-
lsys2untrust match source-address any
set security policies from-zone lsys2-trust to-zone lsys2-untrust policy lsys2trust-to-
lsys2untrust match destination-address any
set security policies from-zone lsys2-trust to-zone lsys2-untrust policy lsys2trust-to-
lsys2untrust match application any
set security policies from-zone lsys2-trust to-zone lsys2-untrust policy lsys2trust-to-
lsys2untrust then permit
set security policies from-zone lsys2-untrust to-zone lsys2-trust policy lsys2untrust-to-
lsys2trust match source-address any
set security policies from-zone lsys2-untrust to-zone lsys2-trust policy lsys2untrust-to-
lsys2trust match destination-address any
set security policies from-zone lsys2-untrust to-zone lsys2-trust policy lsys2untrust-to-
lsys2trust match application any
set security policies from-zone lsys2-untrust to-zone lsys2-trust policy lsys2untrust-to-
lsys2trust then permit
set security policies from-zone lsys2-untrust to-zone lsys2-untrust policy lsys2untrust-to-
lsys2untrust match source-address any
set security policies from-zone lsys2-untrust to-zone lsys2-untrust policy lsys2untrust-to-
lsys2untrust match destination-address any
set security policies from-zone lsys2-untrust to-zone lsys2-untrust policy lsys2untrust-to-
lsys2untrust match application any
set security policies from-zone lsys2-untrust to-zone lsys2-untrust policy lsys2untrust-to-
lsys2untrust then permit
set security policies from-zone lsys2-trust to-zone lsys2-trust policy lsys2trust-to-lsys2trust
match source-address any
set security policies from-zone lsys2-trust to-zone lsys2-trust policy lsys2trust-to-lsys2trust
```

```
match destination-address any
set security policies from-zone lsys2-trust to-zone lsys2-trust policy lsys2trust-to-lsys2trust
match application any
set security policies from-zone lsys2-trust to-zone lsys2-trust policy lsys2trust-to-lsys2trust
then permit
```

Step-by-Step Procedure

NOTE: The user logical system administrator performs the following configuration while logged in to his or her user logical system. The primary administrator can also configure a user logical system at the [edit logical-systems *logical-system*] hierarchy level.

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

To configure the LSYS1 user logical system:

1. Configure interfaces.

```
[edit interfaces]
lsys1-admin@host:LSYS1# set reth1 unit 0 family inet6 address 8885::1/64
lsys1-admin@host:LSYS1# set reth2 unit 0 family inet6 address 7775::1/64
```

2. Configure routing.

```
[edit routing-instances]
lsys1-admin@host:LSYS1# set vr11 instance-type virtual-router
lsys1-admin@host:LSYS1# set vr11 interface lt-0/0/0.3
lsys1-admin@host:LSYS1# set vr11 interface reth1.0
lsys1-admin@host:LSYS1# set vr11 routing-options rib vr11.inet6.0 static route 6665::/64 next-
hop 2111::5
lsys1-admin@host:LSYS1# set vr11 routing-options rib vr11.inet6.0 static route 9995::/64 next-
hop 2111::1
lsys1-admin@host:LSYS1# set vr12 instance-type virtual-router
lsys1-admin@host:LSYS1# set vr12 interface reth2.0
lsys1-admin@host:LSYS1# set vr12 routing-options interface-routes rib-group inet6 vr11vr12v6
lsys1-admin@host:LSYS1# set vr12 routing-options rib vr12.inet6.0 static route 8885::/64 next-
table vr11.inet6.0
lsys1-admin@host:LSYS1# set vr12 routing-options rib vr12.inet6.0 static route 9995::/64 next-
```

```
table vr11.inet6.0
lsys1-admin@host:LSYS1# set vr12 routing-options rib vr12.inet6.0 static route 6665::/64 next-
table vr11.inet6.0
lsys1-admin@host:LSYS1# set vr12 routing-options rib vr12.inet6.0 static route 2111::/64 next-
table vr11.inet6.0
```

```
[edit routing-options]
lsys1-admin@host:LSYS1# set rib-groups vr11vr12v6 import-rib vr11.inet6.0
lsys1-admin@host:LSYS1# set rib-groups vr11vr12v6 import-rib vr12.inet6.0
```

3. Configure zones and security policies.

```
[edit security zones]
lsys1-admin@host:LSYS1# set security-zone lsys1-trust host-inbound-traffic system-services all
lsys1-admin@host:LSYS1# set security-zone lsys1-trust host-inbound-traffic protocols all
lsys1-admin@host:LSYS1# set security-zone lsys1-trust interfaces reth1.0
lsys1-admin@host:LSYS1# set security-zone lsys1-trust interfaces lt-0/0/0.3
lsys1-admin@host:LSYS1# set security-zone lsys1-untrust host-inbound-traffic system-services
all
lsys1-admin@host:LSYS1# set security-zone lsys1-untrust host-inbound-traffic protocols all
lsys1-admin@host:LSYS1# set security-zone lsys1-untrust interfaces reth2.0
```

```
[edit security policies from-zone lsys1-trust to-zone lsys1-untrust]
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1untrust match source-address any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1untrust match destination-address any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1untrust match application any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1untrust then permit
```

```
[edit security policies from-zone lsys1-untrust to-zone lsys1-trust]
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1trust match source-address any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1trust match destination-address any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1trust match application any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1trust then permit
```

```
[edit security policies from-zone lsys1-untrust to-zone lsys1-untrust]
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1untrust match source-address any
```

```
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1untrust match destination-address any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1untrust match application any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1untrust then permit
```

```
[edit security policies from-zone lsys1-trust to-zone lsys1-trust]
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1trust match source-address any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1trust match destination-address any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1trust match application any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1trust then permit
```

Step-by-Step Procedure

To configure the LSYS2 user logical system:

1. Configure interfaces.

```
[edit interfaces]
lsys2-admin@host:LSYS2# set reth3 unit 0 family inet6 address 6665::1/64
```

2. Configure routing.

```
[edit routing-instances]
lsys2-admin@host:LSYS2# set vr2 instance-type virtual-router
lsys2-admin@host:LSYS2# set vr2 interface lt-0/0/0.5
lsys2-admin@host:LSYS2# set vr2 interface reth3.0
lsys2-admin@host:LSYS2# set vr2 routing-options rib vr2.inet6.0 static route 7775::/64 next-
hop 2111::3
lsys2-admin@host:LSYS2# set vr2 routing-options rib vr2.inet6.0 static route 8885::/64 next-
hop 2111::3
lsys2-admin@host:LSYS2# set vr2 routing-options rib vr2.inet6.0 static route 9995::/64 next-
hop 2111::1
```

3. Configure zones and security policies.

```
[edit security zones]
lsys2-admin@host:LSYS2# set security-zone lsys2-trust host-inbound-traffic system-services all
lsys2-admin@host:LSYS2# set security-zone lsys2-trust host-inbound-traffic protocols all
lsys2-admin@host:LSYS2# set security-zone lsys2-trust interfaces reth3.0
```

```
lsys2-admin@host:LSYS2# set security zones security-zone lsys2-untrust host-inbound-traffic
system-services all
lsys2-admin@host:LSYS2# set security-zone lsys2-untrust host-inbound-traffic protocols all
lsys2-admin@host:LSYS2# set security-zone lsys2-untrust interfaces lt-0/0/0.5
```

```
[edit security policies from-zone lsys2-trust to-zone lsys2-untrust]
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2untrust match source-address any
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2untrust match destination-address any
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2untrust match application any
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2untrust then permit
```

```
[edit security policies from-zone from-zone lsys2-untrust to-zone lsys2-trust]
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2trust match source-address any
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2trust match destination-address any
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2trust match application any
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2trust then permit
```

```
[edit security policies from-zone lsys2-untrust to-zone lsys2-untrust]
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2untrust match source-address any
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2untrust match destination-address any
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2untrust match application any
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2untrust then permit
```

```
[edit security policies from-zone lsys2-trust to-zone lsys2-trust]
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2trust match source-address any
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2trust match destination-address any
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2trust match application any
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2trust then permit
```

Results

From configuration mode, confirm the configuration for LSYS1 by entering the `show interfaces`, `show routing-instances`, `show routing-options`, and `show security` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
lsys1-admin@host:LSYS1# show interfaces
interfaces {
    lt-0/0/0 {
        unit 3 {
            encapsulation ethernet;
            peer-unit 2;
            family inet6 {
                address 2111::3/64;
            }
        }
    }
    reth1 {
        unit 0 {
            family inet6 {
                address 8885::1/64;
            }
        }
    }
    reth2 {
        unit 0 {
            family inet6 {
                address 7775::1/64;
            }
        }
    }
}
[edit]
lsys1-admin@host:LSYS1# show routing-instances
routing-instances {
    vr11 {
        instance-type virtual-router;
        interface lt-0/0/0.3;
        interface reth1.0;
        routing-options {
            rib vr11.inet6.0 {
```

```
                static {
                    route 6665::/64 next-hop 2111::5;
                    route 9995::/64 next-hop 2111::1;
                }
            }
        }
    }
    vr12 {
        instance-type virtual-router;
        interface reth2.0;
        routing-options {
            interface-routes {
                rib-group inet6 vr11vr12v6;
            }
            rib vr12.inet6.0 {
                static {
                    route 8885::/64 next-table vr11.inet6.0;
                    route 9995::/64 next-table vr11.inet6.0;
                    route 6665::/64 next-table vr11.inet6.0;
                    route 2111::/64 next-table vr11.inet6.0;
                }
            }
        }
    }
}
[edit]
lsys1-admin@host:LSYS1# show routing-options
rib-groups {
    vr11vr12v6 {
        import-rib [ vr11.inet6.0 vr12.inet6.0 ];
    }
}
[edit]
lsys1-admin@host:LSYS1# show security
security {
    policies {
        from-zone lsys1-trust to-zone lsys1-untrust {
            policy lsys1trust-to-lsys1untrust {
                match {
                    source-address any;
                    destination-address any;
                    application any;
                }
```

```
                then {
                    permit;
                }
            }
        }
        from-zone lsys1-untrust to-zone lsys1-trust {
            policy lsys1untrust-to-lsys1trust {
                match {
                    source-address any;
                    destination-address any;
                    application any;
                }
                then {
                    permit;
                }
            }
        }
        from-zone lsys1-untrust to-zone lsys1-untrust {
            policy lsys1untrust-to-lsys1untrust {
                match {
                    source-address any;
                    destination-address any;
                    application any;
                }
                then {
                    permit;
                }
            }
        }
        from-zone lsys1-trust to-zone lsys1-trust {
            policy lsys1trust-to-lsys1trust {
                match {
                    source-address any;
                    destination-address any;
                    application any;
                }
                then {
                    permit;
                }
            }
        }
    }
    zones {
```

```
        security-zone lsys1-trust {
            host-inbound-traffic {
                system-services {
                    all;
                }
                protocols {
                    all;
                }
            }
            interfaces {
                reth1.0;
                lt-0/0/0.3;
            }
        }
        security-zone lsys1-untrust {
            host-inbound-traffic {
                system-services {
                    all;
                }
                protocols {
                    all;
                }
            }
            interfaces {
                reth2.0;
            }
        }
    }
}
```

From configuration mode, confirm the configuration for LSYS2 by entering the `show interfaces`, `show routing-instances`, and `show security` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
lsys2-admin@host:LSYS2# show interfaces
interfaces {
    lt-0/0/0 {
        unit 5 {
            encapsulation ethernet;
            peer-unit 4;
            family inet6 {
```

```
                address 2111::5/64;
            }
        }
    }
    reth3 {
        unit 0 {
            family inet6 {
                address 6665::1/64;
            }
        }
    }
}
[edit]
lsys2-admin@host:LSYS2# show routing-instances
routing-instances {
    vr2 {
        instance-type virtual-router;
        interface lt-0/0/0.5;
        interface reth3.0;
        routing-options {
            rib vr2.inet6.0 {
                static {
                    route 7775::/64 next-hop 2111::3;
                    route 8885::/64 next-hop 2111::3;
                    route 9995::/64 next-hop 2111::1;
                }
            }
        }
    }
}
[edit]
lsys2-admin@host:LSYS2# show security
security {
    policies {
        from-zone lsys2-trust to-zone lsys2-untrust {
            policy lsys2trust-to-lsys2untrust {
                match {
                    source-address any;
                    destination-address any;
                    application any;
                }
                then {
                    permit;
```

```
                }
            }
        }
        from-zone lsys2-untrust to-zone lsys2-trust {
            policy lsys2untrust-to-lsys2trust {
                match {
                    source-address any;
                    destination-address any;
                    application any;
                }
                then {
                    permit;
                }
            }
        }
        from-zone lsys2-untrust to-zone lsys2-untrust {
            policy lsys2untrust-to-lsys2untrust {
                match {
                    source-address any;
                    destination-address any;
                    application any;
                }
                then {
                    permit;
                }
            }
        }
        from-zone lsys2-trust to-zone lsys2-trust {
            policy lsys2trust-to-lsys2trust {
                match {
                    source-address any;
                    destination-address any;
                    application any;
                }
                then {
                    permit;
                }
            }
        }
    }
    zones {
        security-zone lsys2-trust {
            host-inbound-traffic {
```

```
                system-services {
                    all;
                }
                protocols {
                    all;
                }
            }
            interfaces {
                reth3.0;
            }
        }
        security-zone lsys2-untrust {
            host-inbound-traffic {
                system-services {
                    all;
                }
                protocols {
                    all;
                }
            }
            interfaces {
                lt-0/0/0.5;
            }
        }
    }
}
```

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

Verification

IN THIS SECTION

- Verifying Chassis Cluster Status (IPv6) | **493**
- Troubleshooting Chassis Cluster with Logs (IPv6) | **493**
- Verifying Logical System Licenses (IPv6) | **494**
- Verifying Logical System License Usage (IPv6) | **494**
- Verifying Intra-Logical System Traffic on a Logical System (IPv6) | **495**
- Verifying Intra-Logical System Traffic Within All Logical Systems (IPv6) | **496**

- Verifying Traffic Between User Logical Systems (IPv6) | **498**

Confirm that the configuration is working properly.

Verifying Chassis Cluster Status (IPv6)

Purpose

Verify the chassis cluster status, failover status, and redundancy group information.

Action

From operational mode, enter the `show chassis cluster status` command.

```
{primary:node0}
show chassis cluster status
Cluster ID: 1
Node                  Priority          Status    Preempt  Manual failover

Redundancy group: 0 , Failover count: 1
    node0                   200         primary        no       no
    node1                   100         secondary      no       no

Redundancy group: 1 , Failover count: 1
    node0                   200         primary        no       no
    node1                   100         secondary      no       no
```

Troubleshooting Chassis Cluster with Logs (IPv6)

Purpose

Use these logs to identify any chassis cluster issues. You should run these logs on both nodes.

Action

From operational mode, enter these `show log` commands.

```
user@host> show log jsrpd
user@host> show log chassisd
user@host> show log messages
user@host> show log dcd
user@host> show traceoptions
```

Verifying Logical System Licenses (IPv6)

Purpose

Verify information about logical system licenses.

Action

From operational mode, enter the `show system license status logical-system all` command.

```
{primary:node0}
user@host> show system license status logical-system all
node0:
--------------------------------------------------------------------------
Logical system license status:

logical system name                license status
root-logical-system                enabled
LSYS0                              enabled
LSYS1                              enabled
LSYS2                              enabled
```

Verifying Logical System License Usage (IPv6)

Purpose

Verify information about logical system license usage.

NOTE: The actual number of licenses used is only displayed on the primary node.

Action

From operational mode, enter the `show system license` command.

```
{primary:node0}
user@host> show system license
License usage:
                                Licenses     Licenses    Licenses    Expiry
  Feature name                       used    installed      needed
  logical-system                        4           25           0    permanent

Licenses installed:
  License identifier: JUNOS305013
  License version: 2
  Valid for device: JN110B54BAGB
  Features:
    logical-system-25 - Logical System Capacity
      permanent
```

Verifying Intra-Logical System Traffic on a Logical System (IPv6)

Purpose

Verify information about currently active security sessions within a logical system.

Action

From operational mode, enter the `show security flow session logical-system LSYS1` command.

```
{primary:node0}
user@host> show security flow session logical-system LSYS1
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
```

```
Session ID: 10000115, Policy name: lsys1trust-to-lsys1untrust/8, State: Active, Timeout: 1784,
Valid
  In: 8885::2/34564 --> 7775::2/23;tcp, If: reth1.0, Pkts: 22, Bytes: 1745
  Out: 7775::2/23 --> 8885::2/34564;tcp, If: reth2.0, Pkts: 19, Bytes: 2108
Total sessions: 1

Flow Sessions on FPC2 PIC0:
Total sessions: 0

Flow Sessions on FPC2 PIC1:
Total sessions: 0

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

Flow Sessions on FPC0 PIC1:

Session ID: 10000006, Policy name: lsys1trust-to-lsys1untrust/8, State: Backup, Timeout: 14392,
Valid
  In: 8885::2/34564 --> 7775::2/23;tcp, If: reth1.0, Pkts: 0, Bytes: 0
  Out: 7775::2/23 --> 8885::2/34564;tcp, If: reth2.0, Pkts: 0, Bytes: 0
Total sessions: 1

Flow Sessions on FPC2 PIC0:
Total sessions: 0

Flow Sessions on FPC2 PIC1:
Total sessions: 0
```

Verifying Intra-Logical System Traffic Within All Logical Systems (IPv6)

Purpose

Verify information about currently active security sessions on all logical systems.

Action

From operational mode, enter the `show security flow session logical-system all` command.

```
{primary:node0}
user@host> show security flow session logical-system all
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000115, Policy name: lsys1trust-to-lsys1untrust/8, State: Active, Timeout: 1776,
Valid
Logical system: LSYS1
  In: 8885::2/34564 --> 7775::2/23;tcp, If: reth1.0, Pkts: 22, Bytes: 1745
  Out: 7775::2/23 --> 8885::2/34564;tcp, If: reth2.0, Pkts: 19, Bytes: 2108
Total sessions: 1

Flow Sessions on FPC2 PIC0:
Total sessions: 0

Flow Sessions on FPC2 PIC1:
Total sessions: 0

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

Flow Sessions on FPC0 PIC1:

Session ID: 10000006, Policy name: lsys1trust-to-lsys1untrust/8, State: Backup, Timeout: 14384,
Valid
Logical system: LSYS1
  In: 8885::2/34564 --> 7775::2/23;tcp, If: reth1.0, Pkts: 0, Bytes: 0
  Out: 7775::2/23 --> 8885::2/34564;tcp, If: reth2.0, Pkts: 0, Bytes: 0
Total sessions: 1

Flow Sessions on FPC2 PIC0:
Total sessions: 0

Flow Sessions on FPC2 PIC1:
Total sessions: 0
```

Verifying Traffic Between User Logical Systems (IPv6)

Purpose

Verify information about currently active security sessions between logical systems.

Action

From operational mode, enter the `show security flow session logical-system` *logical-system-name* command.

```
{primary:node0}
user@host> show security flow session logical-system LSYS1

node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:

Session ID: 80000118, Policy name: lsys1trust-to-lsys1trust/11, State: Active, Timeout: 1792,
Valid
  In: 8885::2/34565 --> 6665::2/23;tcp, If: reth1.0, Pkts: 91, Bytes: 6802
  Out: 6665::2/23 --> 8885::2/34565;tcp, If: lt-0/0/0.3, Pkts: 65, Bytes: 6701
Total sessions: 1

Flow Sessions on FPC2 PIC1:
Total sessions: 0

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

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:

Session ID: 80000010, Policy name: lsys1trust-to-lsys1trust/11, State: Backup, Timeout: 14388,
Valid
  In: 8885::2/34565 --> 6665::2/23;tcp, If: reth1.0, Pkts: 0, Bytes: 0
  Out: 6665::2/23 --> 8885::2/34565;tcp, If: lt-0/0/0.3, Pkts: 0, Bytes: 0
```

```
Total sessions: 1

Flow Sessions on FPC2 PIC1:
Total sessions: 0
```

```
{primary:node0}
user@host> show security flow session logical-system LSYS2

node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:

Session ID: 80000119, Policy name: lsys2untrust-to-lsys2trust/13, State: Active, Timeout: 1788,
Valid
  In: 8885::2/34565 --> 6665::2/23;tcp, If: lt-0/0/0.5, Pkts: 91, Bytes: 6802
  Out: 6665::2/23 --> 8885::2/34565;tcp, If: reth3.0, Pkts: 65, Bytes: 6701
Total sessions: 1

Flow Sessions on FPC2 PIC1:
Total sessions: 0

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

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:

Session ID: 80000011, Policy name: lsys2untrust-to-lsys2trust/13, State: Backup, Timeout: 14380,
Valid
  In: 8885::2/34565 --> 6665::2/23;tcp, If: lt-0/0/0.5, Pkts: 0, Bytes: 0
  Out: 6665::2/23 --> 8885::2/34565;tcp, If: reth3.0, Pkts: 0, Bytes: 0
Total sessions: 1
```

```
Flow Sessions on FPC2 PIC1:
Total sessions: 0
```

```
{primary:node0}
user@host> show security flow session logical-system all

node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:

Session ID: 80000118, Policy name: lsys1trust-to-lsys1trust/11, State: Active, Timeout: 1784,
Valid
Logical system: LSYS1
  In: 8885::2/34565 --> 6665::2/23;tcp, If: reth1.0, Pkts: 91, Bytes: 6802
  Out: 6665::2/23 --> 8885::2/34565;tcp, If: lt-0/0/0.3, Pkts: 65, Bytes: 6701

Session ID: 80000119, Policy name: lsys2untrust-to-lsys2trust/13, State: Active, Timeout: 1784,
Valid
Logical system: LSYS2
  In: 8885::2/34565 --> 6665::2/23;tcp, If: lt-0/0/0.5, Pkts: 91, Bytes: 6802
  Out: 6665::2/23 --> 8885::2/34565;tcp, If: reth3.0, Pkts: 65, Bytes: 6701
Total sessions: 2

Flow Sessions on FPC2 PIC1:
Total sessions: 0

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

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:

Session ID: 80000010, Policy name: lsys1trust-to-lsys1trust/11, State: Backup, Timeout: 14378,
Valid
Logical system: LSYS1
```

```
  In: 8885::2/34565 --> 6665::2/23;tcp, If: reth1.0, Pkts: 0, Bytes: 0
  Out: 6665::2/23 --> 8885::2/34565;tcp, If: lt-0/0/0.3, Pkts: 0, Bytes: 0

Session ID: 80000011, Policy name: lsys2untrust-to-lsys2trust/13, State: Backup, Timeout: 14376,
Valid
Logical system: LSYS2
  In: 8885::2/34565 --> 6665::2/23;tcp, If: lt-0/0/0.5, Pkts: 0, Bytes: 0
  Out: 6665::2/23 --> 8885::2/34565;tcp, If: reth3.0, Pkts: 0, Bytes: 0
Total sessions: 2

Flow Sessions on FPC2 PIC1:
Total sessions: 0
```

SEE ALSO

- Understanding Logical Systems in the Context of Chassis Cluster | **410**
- Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (Primary Administrators Only) | **411**
- Example: Configuring an Active/Passive Chassis Cluster on SRX5800 Devices
- Chassis Cluster Overview

Release History Table

Release	Description
12.3X48-D50	Starting with Junos OS Release 12.3X48-D50, when you configure the logical systems within a chassis cluster, if logical systems licenses on backup node are not sufficient when you `commit` the configuration, a warning message is displayed about the number of licenses required on backup node as well, just as on primary node in all the previous releases.

Flow Trace for Logical Systems

IN THIS SECTION

- Flow Trace Support for Logical Systems Overview | **502**
- Configure Flow Trace Support for Logical Systems | **502**

Flow trace also called traceoptions, allows you to monitor traffic flow into and out of an SRX Series device. You can use tracoptions as debugging tool to trace the packets as they traverse the SRX Series device. Traceoptions help you to get details of actions by your security device.

Flow Trace Support for Logical Systems Overview

For an SRX Series device configured with logical systems, by default the traceoptions are configured at the root level only. In this case, all the system traces including root and logical systems are logged in one single trace file. This generated large amounts of information in a single file.

Starting in Junos OS Release 19.4R1, you can enable tracing operations per logical system level. When you configure the traceoptions at the logical system level, then the traces for that specific logical systems are logged in the respective trace file. You can generate an output file for the specified logical system, and you can find the required traffic information easily in the trace file.

When you enable traceoptions, you specify the name of the file and the type of information you want to trace.

All flow trace sent to one log file in root, if you enable the traceoptions under root context. Traces for a logical system only sent to the respective trace file, if you enable the traceoptions for the specific logical system.

Configure Flow Trace Support for Logical Systems

Configuring traceoptions for a logical system includes configuring both a target file and a flag. The target file determines where the trace output is recorded. The flag defines what type of data to be collected. If you configure traceoptions for a logical system, the respective trace file sent to the specific logical system log file only.

To configure traceoptions for a logical system:

1. Create logical system LSYS1 and setup the basic configurations. See "Setting Up a Logical System" on page 55
2. Configure target file to save the trace information for the logical system.

```
[edit]
user@host# set logical-systems LSYS1 security flow traceoptions file flow_lsys1.log
user@host# set logical-systems LSYS1 security flow traceoptions file size 1g
```

3. Configure traceoptions flag for the logical system.

```
[edit]
user@host# set logical-systems LSYS1 security flow traceoptions flag all
```

After you commit the traceoptions configuration, you can view the traceoptions debug files for the logical system using `show log` *`tracefilename`* operational command.

```
user@host:LSYS1> show log flow_lsys1.log
Nov  7 07:34:09 07:34:09.491800:CID-0:THREAD_ID-00:LSYS_ID-01:RT:got route table lock

Nov  7 07:34:09 07:34:09.491809:CID-0:THREAD_ID-00:LSYS_ID-01:RT:released route table lock

Nov  7 07:34:09 07:34:09.491840:CID-0:THREAD_ID-00:LSYS_ID-01:RT:got route table lock

Nov  7 07:34:09 07:34:09.491841:CID-0:THREAD_ID-00:LSYS_ID-01:RT:released route table lock

Nov  7 07:34:09 07:34:09.491854:CID-0:THREAD_ID-00:LSYS_ID-01:RT:cache final sw_nh 0x0

Nov  7 07:34:09 07:34:09.491868:CID-0:THREAD_ID-00:LSYS_ID-01:RT:got route table lock

Nov  7 07:34:09 07:34:09.491869:CID-0:THREAD_ID-00:LSYS_ID-01:RT:released route table lock

Nov  7 07:34:09 07:34:09.491881:CID-0:THREAD_ID-00:LSYS_ID-01:RT:cache final sw_nh 0x0
```

Example: Deleting a Logical System

IN THIS SECTION

- Requirements | 504
- Overview | 504
- Configuration | 505
- Verification | 508

This example shows how to delete a logical system configured for an SRX Series Services Gateway device running logical systems. Only the primary administrator can delete a logical system.

Requirements

The example uses an SRX5600 device running Junos OS with Logical Systems.

Alternatively, follow those instructions substituting your own configuration values.

Overview

IN THIS SECTION

- Topology | 504

This example shows how to delete a logical system, which you can do at any time. However, if you have configured the device to include the maximum number of logical systems that are supported you must first delete an existing logical system before you can add another one.

Deletion of a logical system is a simple procedure that includes these tasks:

- Remove from the logical system the security profile that is bound to it.

  Note that in this step you are not deleting the security profile–it might be used for other logical systems–but simply detaching it from the logical system that you intend to delete.

- Detach from the logical system any login classes that are associated with it.

  Removing them from the logical system does not delete the login classes.

- Delete the logical system.

Topology

Configuration

IN THIS SECTION

- Procedure | 505

Procedure

CLI Quick Configuration

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

```
delete system security-profile  ls-design-profile logical-system ls-product-design
delete system login class ls-design-admin logical-system  ls-product-design
delete system login class ls-design-user logical-system  ls-product-design
delete logical-system ls-product-design
```

Step-by-Step Procedure

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

To delete a logical system:

1. Determine that the logical system that you want to delete exists.

```
[edit]
user@host# show logical-systems ?
interconnect-logical-system  Logical system name
ls-accounting-dept   Logical system name
ls-marketing-dept    Logical system name
ls-product-design    Logical system name
```

2. Delete the security profile.

Step-by-Step Procedure

a. Verify that security profile that you intend to detach from the logical system is bound to it.

```
[edit]
user@host# show system security-profile ls-design-profile
logical-system [ ls-product-design ];
```

b. Detach the security profile from the logical system.

```
[edit]
user@host# delete system security-profile  ls-design-profile logical-system ls-product-
design
```

3. Delete the login classes.

Step-by-Step Procedure

a. Display the login class and login user configurations for the user logical system administrator.

```
user@host> show configuration system login class ls-design-admin
logical-system ls-product-design;
permissions all;
user@host> show configuration system login user lsdesignadmin1
full-name lsdesignadmin1;
 uid 2006;
class ls-design-admin;
authentication {
    encrypted-password "$ABC123"; ## SECRET-DATA
    }
```

b. Detach the login class for the administrator from the logical system.

```
[edit]
user@host# delete system login class ls-design-admin logical-system  ls-product-design
```

c. Display the login class and login user configurations for the user.

```
user@host> show configuration system login class ls-design-user
logical-system ls-product-design;
permissions view;
user@host> show configuration system login user lsdesignuser1
full-name lsdesignuser1
 uid 2007;
class ls-design-user;
authentication {
    encrypted-password "$ABC123"; ## SECRET-DATA
    }
```

d. Detach the login class for the user from the logical system.

```
user@host# delete system login class ls-design-user logical-system  ls-product-design
```

4. Delete the logical system.

```
[edit]
user@host# delete logical-system ls-product-design
```

Results

From configuration mode, confirm your configuration by entering the `show logical-systems` command. In this case, the logical system that you deleted should not be included in displayed list of logical systems configured for the device. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show logical-systems
 interconnect-logical-system  Logical system name
ls-accounting-dept   Logical system name
 interconnect-logical-system  Logical system name
ls-marketing-dept    Logical system name
```

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

Verification

IN THIS SECTION

- Verifying That the Correct Logical System and Its Profile and Attached Class Were Deleted | 508

To confirm that the configuration is working properly, perform these tasks:

Verifying That the Correct Logical System and Its Profile and Attached Class Were Deleted

Purpose

Verify if the logical system has been deleted using the show command described previously.

RELATED DOCUMENTATION

Understanding User Logical Systems and the User Logical System Administrator Role | 53

Understanding Logical Systems for SRX Series Services Gateways | 5

Troubleshooting Logical Systems

IN THIS SECTION

- Understanding Security Logs and Logical Systems | 509
- Configuring On-Box Reporting for logical Systems | 511
- Example: Configure Security Log for Logical Systems | 512
- Configuring On-Box Binary Security Log Files for Logical System | 517
- Configuring Off-Box Binary Security Log Files for Logical System | 519
- Understanding Data Path Debugging for Logical Systems | 520

- Performing Tracing for Logical Systems (Primary Administrators Only) | **521**
- Troubleshooting DNS Name Resolution in Logical System Security Policies (Primary Administrators Only) | **527**

Use the following features to monitor logical systems and troubleshoot the software issues. For more information, see the following topics:

Understanding Security Logs and Logical Systems

IN THIS SECTION

- Limitations | **511**

Security logs are system log messages that include security events. If a device is configured for logical systems, security logs generated within the context of a logical system use the name ***logname*_LS** (for example, **IDP_ATTACK_LOG_EVENT_LS**). The logical system version of a log has the same set of attributes as the log for devices that are not configured for logical systems. The logical system log includes logical-system-name as the first attribute.

The following security log shows the attributes for the IDP_ATTACK_LOG_EVENT log for a device that is *not* configured for logical systems:

```
IDP_ATTACK_LOG_EVENT {
help "IDP attack log";
description "IDP Attack log generated for attack";
type event;
args timestamp message-type source-address source-port destination-address destination-port
protocol-name service-name application-name rule-name rulebase-name policy-name repeat-count
action threat-severity attack-name nat-source-address nat-source-port nat-destination-address
nat-destination-port elapsed-time inbound-bytes outbound-bytes inbound-packets outbound-packets
source-zone-name source-interface-name destination-zone-name destination-interface-name packet-
log-id message;
severity LOG_INFO;
flag auditable;
```

```
edit "2010/10/01 mvr created";
}
```

The following security log shows the attributes for the IDP_ATTACK_LOG_EVENT_LS log for a device that is configured for logical systems (note that logical-system-name is the first attribute):

```
IDP_ATTACK_LOG_EVENT_LS {
help "IDP attack log";
description "IDP Attack log generated for attack";
type event;
args logical-system-name timestamp message-type source-address source-port destination-address
destination-port protocol-name service-name application-name rule-name rulebase-name policy-name
repeat-count action threat-severity attack-name nat-source-address nat-source-port nat-
destination-address nat-destination-port elapsed-time inbound-bytes outbound-bytes inbound-
packets outbound-packets source-zone-name source-interface-name destination-zone-name
destination-interface-name packet-log-id message;
severity LOG_INFO;
flag auditable;
edit "2010/10/01 mvr created";
}
```

If a device is configured for logical systems, log parsing scripts might need to be modified because the log name includes the **_LS** suffix and the logical-system-name attribute can be used to segregate logs by logical system.

If a device is not configured for logical systems, the security logs remain unchanged and scripts built to parse logs do not need any modification.

NOTE: Only the primary administrator can configure logging at the [`edit security log`] hierarchy level. User logical system administrators cannot configure logging for their logical systems.

Stream mode is a set of logging services that includes:

- Off-box logging (SRX Series)
- On-box logging and reporting (SRX300, SRX320, SRX340, SRX345, SRX380, SRX550M, SRX1500, SRX4100, SRX4200, and SRX4600 Series)

Per logical system configuration is supported for the off-box logging and logs are handled based on these configurations. Previously the user logical system logs were generated from root logical system. For off-box logging, the logical system logs can only be generated from logical system interface.

Limitations

Each SPU can only support a maximum of 1000 connections for standalone and 500 connections for cluster on the SRX5400, SRX5600, and SRX5800 devices in the Junos OS 18.2R1 release. If all the connections are used up, some connections for user logical systems might not be established.

NOTE: The error message will be captured in the System Log Explorer.

Configuring On-Box Reporting for logical Systems

SRX Series devices supports different types of reports for logical system users.

Reports are stored locally on the SRX Series device and there is no requirement for separate devices or tools for logs and reports storage. The on-box reports provides a simple and easy-to-use interface for viewing the security logs.

Before you begin:

- Understand how to configure security log for logical systems. See Example: Configure Security Log for logical Systems

To configure on-box reporting for logical system:

1. Define the logical system name as LSYS1.

```
user@host# set logical-systems LSYS1
```

2. Create report within security log per tenant system.

```
user@host# set logical-systems LSYS1 security log report
```

3. Confirm your configuration by entering the `show logical-systems LSYS1` command.

```
user@host# show logical-systems LSYS1
security {
    log {
        report;
```

```
    }
}
```

NOTE: By default the `report` option is disabled. The `set logical-systems LSYS1 security log mode stream` command is enabled by default.

Example: Configure Security Log for Logical Systems

IN THIS SECTION

- Requirements | **512**
- Overview | **512**
- Configuration | **513**
- Verification | **516**

This example shows how to configure security logs for a logical system.

Requirements

This example uses the following hardware and software components:

- An SRX Series device.
- Junos OS Release 18.3R1 and later releases.

Before you begin:

- Understand how to configure a logical system.
- Understand how to create security profiles for the primary logical system. See Understanding Logical Systems Security Profiles (Primary Administrators Only).

Overview

SRX Series devices have two types of log: system logs and security logs. System logs record control plane events, for example, admin login to the device. Security logs, also known as traffic logs, record

data plane events regarding specific traffic handling, for example when a security policy denies certain traffic due to some violation of the policy.

The two types of logs can be collected and saved either on-box or off-box. The procedure below explains how to configure security logs in binary format for off-box (stream-mode) logging.

For off-box logging, security logs for a logical system are sent from a logical system interface. If the logical system interface is already configured in a routing instance, then configure `routing-instance routing-instance-name` at `edit logical-systems logical-system-name security log stream log-stream-name host` hierarchy. If the interface is not configured in routing instance, then no routing instance should be configured at `edit logical-systems logical-system-name security log stream log-stream-name host` hierarchy.

Configuration

IN THIS SECTION

- CLI Quick Configuration | **513**
- Procedure | **514**
- Procedure | **515**
- Results | **515**

CLI Quick Configuration

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

```
set logical-systems LSYS1 security log mode stream
set logical-systems LSYS1 security log stream LSYS1_s format binary host 1.3.54.22
set logical-systems LSYS1 security log source-address 2.3.45.66
set logical-systems LSYS1 security log transport protocol tls
set logical-systems LSYS1 routing-instances LSYS1_ri instance-type virtual-router
set logical-systems LSYS1 routing-instances LSYS_ri interface ge-0/0/3
set logical-systems LSYS1 security log stream LSYS1_s host routing-instance LSYS1_ri
set system security-profile p1 security-log-stream-number reserved 1
set system security-profile p1 security-log-stream-number maximum 2
set system security-profile LSYS1_profile logical-system LSYS1
```

Procedure

Step-by-Step Procedure

The following procedure specifies how to configure security logs for a logical system.

1. Specify the logging mode and the format for the log file. For off-box, stream-mode logging.

```
[edit ]
user@host# set logical-systems LSYS1 security log mode stream
user@host# set logical-systems LSYS1 security log stream LSYS1_s format binary host 1.3.54.22
```

2. For off-box security logging, specify the source address, which identifies the SRX Series device that generated the log messages. The source address is required.

```
[edit ]
user@host# set logical-systems LSYS1 security log source-address 2.3.45.66
```

3. Specify the routing instance and define the interface.

```
[edit ]
user@host# set logical-systems LSYS1 routing-instances LSYS1_ri instance-type virtual-router
user@host# set logical-systems LSYS1 routing-instances LSYS_ri interface ge-0/0/3
```

4. Define routing instance for a logical system.

```
[edit ]
user@host# set logical-systems LSYS1 security log stream LSYS1_s host routing-instance
LSYS1_ri
```

5. Specify the security log transport protocol for the device.

```
[edit ]
user@host# set logical-systems LSYS1 security log transport protocol tls
```

Procedure

Step-by-Step Procedure

The following procedure specifies how to configure a security profile for a logical system.

1. Configure a security profile and specify the number of maximum and reserved policies.

```
[edit ]
user@host# set system security-profile p1 security-log-stream-number reserved 1
user@host# set system security-profile p1 security-log-stream-number maximum 2
```

2. Assign the configured security profile to TSYS1.

```
[edit ]
user@host# set system security-profile LSYS1_profile logical-system LSYS1
```

Results

From configuration mode, confirm your configuration by entering the `show system security-profile`, `show logical-systems LSYS1 security log`, and `show logical-systems LSYS1 routing-instances` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show system security-profile
LSYS1_profile {
    logical-system LSYS1;
}
p1 {
    security-log-stream-number {
        maximum 2;
        reserved 1;
    }
}
```

```
[edit]
user@host# show logical-systems LSYS1 security log
```

```
mode stream;
source-address 2.3.45.66;
transport {
    protocol tls;
}
stream LSYS1_s {
    format binary;
    host {
        1.3.54.22;
    }
}
```

```
[edit]
user@host# show logical-systems LSYS1 routing-instances
LSYS1_ri {
    instance-type virtual-router;
    interface ge-0/0/3.0;
}
```

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

Verification

IN THIS SECTION

- Verifying Detailed Output for Security Log | **516**

Verifying Detailed Output for Security Log

Purpose

Verify that the output displays the resource information for all logical systems.

Action

From operational mode, enter the `show system security-profile security-log-stream-number tenant all` command.

```
logical-system name  security profile name  usage    reserved    maximum

root-logical-system         Default-Profile           0          0
```

Meaning

The output displays the resource information for logical systems.

Configuring On-Box Binary Security Log Files for Logical System

SRX Series devices support two types of log: system logs and security logs.

The two types of log are collected and saved either on-box or off-box. The following procedure explains how to configure security logs in binary format for on-box (event-mode) logging for logical system.

The following procedure specifies binary format for event-mode security logging, and defines the log filename, path, and log file characteristics for logical system.

1. Specify the logging mode and the format for the log file. For on-box, event-mode logging:

```
[edit]
user@host# set logical-systems LSYS1 security log mode event
user@host# set logical-systems LSYS1 security log format binary
```

2. (Optional) Specify a log filename.

```
[edit]
user@host# set logical-systems LSYS1 security log file name security-binary-log
```

NOTE: Security log filename is not mandatory. If security log filename is not configured, by default the file bin_messages is created in the /var/log directory.

3. Confirm your configuration by entering the `show logical-systems LSYS1` command.

```
[edit]
user@host# show logical-systems LSYS1
security {
    log {
        mode event;
        format binary;
        file {
            name security-binary-log;
        }
    }
}
```

The following procedure specifies binary format for stream-mode security logging, and defines the log filename and log file characteristics for logical system.

1. Specify the logging mode and the format for the log file. For on-box, stream-mode logging:

```
[edit]
user@host# set logical-systems LSYS1 security log mode stream
user@host# set logical-systems LSYS1 security log stream s1 format binary
```

2. (Optional) Specify a log filename.

```
[edit]
user@host# set logical-systems LSYS1 security log stream s1 file name f1.bin
```

3. Confirm your configuration by entering the `show logical-systems LSYS1` command.

```
[edit]
user@host# show logical-systems LSYS1
security {
    log {
        mode stream;
```

```
        stream s1 {
            format binary;
            file {
                name f1.bin;
            }
        }
    }
}
```

Configuring Off-Box Binary Security Log Files for Logical System

SRX Series devices support two types of log: system logs and security logs.

The two types of log can be collected and saved either on-box or off-box. The procedure below explains how to configure security logs in binary format for off-box (stream-mode) logging.

The following procedure specifies binary format for stream-mode security logging, and defines the logging mode, source address, and host name characteristics for logical system.

1. Specify the logging mode and the format for the log file. For off-box, stream-mode logging:

```
[edit]
user@host# set logical-systems LSYS1 security log mode stream s1 format binary
```

2. Specify the source address for off-box security logging.

```
[edit]
user@host# set logical-systems LSYS1 security log source-address 100.0.0.1
```

3. Specify the host name.

```
[edit]
user@host# set logical-systems LSYS1 security log stream s1 host 100.0.0.2
```

4. Confirm your configuration by entering the `show logical-systems LSYS1` command.

```
[edit]
user@host#show logical-systems LSYS1
security {
    log {
        mode stream;
        source-address 100.0.0.1;
        stream s1 {
            format binary;
            host {
                100.0.0.2;
            }
        }
    }
}
```

Understanding Data Path Debugging for Logical Systems

Data path debugging provides tracing and debugging at multiple processing units along the packet-processing path. Data path debugging can also be performed on traffic between logical systems.

> **NOTE**: Only the primary administrator can configure data path debugging for logical systems at the [**edit security datapath-debug**] level. User logical system administrators cannot configure data path debugging for their logical systems.

End-to-end event tracing traces the path of a packet from when it enters the device to when it leaves the device. When the primary administrator configures end-to-end event tracing, the trace output contains logical system information.

The primary administrator can also configure tracing for traffic between logical systems. The trace output shows traffic entering and leaving the logical tunnel between logical systems. When the **preserve-trace-order** option is configured, the trace message is sorted chronologically. In addition to the trace action, other actions such as packet-dump and packet-summary may be configured for traffic between logical systems.

Data path debugging is supported on SRX1400, SRX3400, SRX3600, SRX5400, SRX5600, and SRX5800.

SEE ALSO

Performing Tracing for Logical Systems (Primary Administrators Only) | 521

Performing Tracing for Logical Systems (Primary Administrators Only)

NOTE: Only the primary administrator can configure data path debugging for logical systems at the root level.

To configure an action profile for a trace or packet capture:

1. Specify event types and trace actions. You can specify any combination of event types and trace actions. For example, the following statements configure multiple trace actions for each event type:

```
[edit security datapath-debug]
user@host# set action-profile p1 event lbt trace
user@host# set action-profile p1 event lbt count
user@host# set action-profile p1 event lbt packet-summary
user@host# set action-profile p1 event lbt packet-dump
user@host# set action-profile p1 event pot trace
user@host# set action-profile p1 event pot count
user@host# set action-profile p1 event pot packet-summary
user@host# set action-profile p1 event pot packet-dump
user@host# set action-profile p1 event np-ingress trace
user@host# set action-profile p1 event np-ingress count
user@host# set action-profile p1 event np-ingress packet-summary
user@host# set action-profile p1 event np-ingress packet-dump
user@host# set action-profile p1 event np-egress trace
user@host# set action-profile p1 event np-egress count
user@host# set action-profile p1 event np-egress packet-summary
user@host# set action-profile p1 event np-egress packet-dump
user@host# set action-profile p1 event jexec trace
user@host# set action-profile p1 event jexec count
user@host# set action-profile p1 event jexec packet-summary
user@host# set action-profile p1 event jexec packet-dump
user@host# set action-profile p1 event lt-enter trace
user@host# set action-profile p1 event lt-enter count
user@host# set action-profile p1 event lt-enter packet-summary
user@host# set action-profile p1 event lt-enter packet-dump
```

```
user@host# set action-profile p1 event lt-leave trace
user@host# set action-profile p1 event lt-leave count
user@host# set action-profile p1 event lt-leave packet-summary
user@host# set action-profile p1 event lt-leave packet-dump
```

2. Specify action profile options.

```
[edit security datapath-debug]
user@host# set action-profile p1 record-pic-history
user@host# set action-profile p1 preserve-trace-order
```

3. Configure packet filter options.

```
[edit security datapath-debug]
user@host# set packet-filter 1 action-profile p1
user@host# set packet-filter 1 protocol udp
```

To capture trace messages for logical systems:

1. Configure the trace capture file.

```
[edit security datapath-debug]
user@host# set traceoptions file e2e.trace
user@host# set traceoptions file size 10m
```

2. Display the captured trace in operational mode.

```
user@host> show log e2e.trace
Jul  7 09:49:56 09:49:56.417578:CID-00:FPC-01:PIC-00:THREAD_ID-00:FINDEX:0:IIF:75:SEQ:0:TC:0
PIC History: ->C0/F1/P0
NP ingress channel 0 packet
Meta: Src: F1/P0 Dst: F0/P0
IP: saddr 10.1.1.2 daddr 30.1.1.2 proto 6 len 500

Jul  7 09:49:56 09:49:55.1414031:CID-00:FPC-00:PIC-00:THREAD_ID-04:FINDEX:0:IIF:75:SEQ:0:TC:1
PIC History: ->C0/F1/P0->C0/F0/P0
LBT pkt, payload: DATA
Meta: Src: F1/P0 Dst: F0/P0
IP: saddr 10.1.1.2 daddr 30.1.1.2 proto 6 len 500
```

```
...

(Some trace information omitted)
...

.Jul  7 09:49:56
09:49:55.1415649:CID-00:FPC-00:PIC-00:THREAD_ID-05:FINDEX:0:IIF:75:SEQ:0:TC:16
PIC History: ->C0/F1/P0->C0/F0/P0->C0/F0/P0->C0/F0/P0->C0/F0/P0
POT pkt, action: POT_SEND payload: DATA
Meta: Src: F0/P0 Dst: F1/P0
IP: saddr 10.1.1.2 daddr 30.1.1.2 proto 6 len 500

Jul  7 09:49:56 09:49:56.419274:CID-00:FPC-01:PIC-00:THREAD_ID-00:FINDEX:0:IIF:75:SEQ:0:TC:17
PIC History: ->C0/F1/P0->C0/F0/P0->C0/F0/P0->C0/F0/P0->C0/F0/P0->C0/F1/P0
NP egress channel 0 packet
Meta: Src: F0/P0 Dst: F1/P0
IP: saddr 10.1.1.2 daddr 30.1.1.2 proto 6 len 500
```

3. Clear the log.

```
user@host> clear log e2e.trace
```

To perform packet capture for logical systems:

1. Configure the packet capture file.

```
[edit security datapath-debug]
user@host# set capture-file e2e.pcap
user@host# set capture-file format pcap
user@host# set capture-file size 10m
user@host# set capture-file world-readable
user@host# set capture-file maximum-capture-size 1500
```

2. Enter operational mode to start and then stop the packet capture.

```
user@host> request security datapath-debug capture start
user@host> request security datapath-debug capture stop
```

NOTE: Packet capture files can be opened and analyzed offline with tcpdump or any packet analyzer that recognizes the libpcap format. You can also use FTP or the Session Control Protocol (SCP) to transfer the packet capture files to an external device.

3. Disable packet capture from configuration mode.

NOTE: Disable packet capture before opening the file for analysis or transferring the file to an external device with FTP or SCP. Disabling packet capture ensures that the internal file buffer is flushed and all the captured packets are written to the file.

```
[edit forwarding-options]
user@host# set packet-capture disable
```

4. Display the packet capture.

- To display the packet capture with the tcpdump utility:

```
user@host# tcpdump -nr /var/log/e2e.pcap
09:49:55.1413990 C0/F0/P0 event:11(lbt) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345: S
0:460(460) win 0
09:49:55.1414154 C0/F0/P0 event:11(lbt) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345: S
0:460(460) win 0
09:49:55.1415062 C0/F0/P0 event:11(lbt) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345: S
0:460(460) win 0
09:49:55.1415184 C0/F0/P0 event:11(lbt) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345: S
0:460(460) win 0
09:49:55.1414093 C0/F0/P0 event:12(pot) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345: S
0:460(460) win 0
09:49:55.1414638 C0/F0/P0 event:12(pot) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345: S
0:460(460) win 0
09:49:55.1415011 C0/F0/P0 event:12(pot) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345: S
0:460(460) win 0
09:49:55.1415129 C0/F0/P0 event:12(pot) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345: S
0:460(460) win 0
09:49:55.1415511 C0/F0/P0 event:12(pot) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345: S
0:460(460) win 0
09:49:55.1415649 C0/F0/P0 event:12(pot) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345: S
0:460(460) win 0
```

```
09:49:55.1415249 C0/F0/P0 event:18(jexec) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345: S
0:460(460) win 0
09:49:55.1415558 C0/F0/P0 event:18(jexec) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345: S
0:460(460) win 0
09:49:55.1414226 C0/F0/P0 event:18(jexec) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345: S
0:460(460) win 0
09:49:55.1414696 C0/F0/P0 event:18(jexec) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345: S
0:460(460) win 0
09:49:55.1414828 C0/F0/P0 event:16(lt-enter) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345: S
0:460(460) win 0
09:49:55.1414919 C0/F0/P0 event:15(lt-leave) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345: S
0:460(460) win 0
09:49:56.417560 C0/F1/P0 event:1(np-ingress) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345: S
0:460(460) win 0
09:49:56.419263 C0/F1/P0 event:2(np-egress) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345: S
0:460(460) win 0
```

- To display the packet capture from CLI operational mode:

```
user@host> show security datapath-debug capture
Packet 1, len 568: (C0/F0/P0/SEQ:0:lbt)
00 00 00 00 00 00 50 c5 8d 0c 99 4a 00 00 0a 01
01 02 08 00 45 60 01 f4 00 00 00 00 40 06 4e 9f
0a 01 01 02 1e 01 01 02 5b 9b 30 39 00 00 00 00
00 00 00 00 50 02 00 00 f8 3c 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 ac 7a 00 04
00 00 00 00 b3 e3 15 4e 66 93 15 00 04 22 38 02
38 02 00 00 00 01 00 03 0b 00 00 00 50 d0 1a 08
30 de be bf e4 f3 19 08
Packet 2, len 624: (C0/F0/P0/SEQ:0:lbt)
aa 35 00 00 00 00 00 00 00 00 00 00 03 00 00
00 0a 00 00 00 00 00 00 05 bd 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 50 c5
8d 0c 99 4a 00 00 0a 01 01 02 08 00 45 60 01 f4
00 00 00 00 40 06 4e 9f 0a 01 01 02 ac 7a 00 04
00 00 00 00 b3 e3 15 4e 0a 94 15 00 04 5a 70 02
70 02 00 00 00 03 00 03 0b 00 00 00 50 d0 1a 08
30 de be bf e4 f3 19 08

...
```

```
(Packets 3 through 17 omitted)
...

Packet 18, len 568: (C0/F1/P0/SEQ:0:np-egress)
00 00 00 04 00 00 00 00 1e 01 01 02 50 c5 8d 0c
99 4b 08 00 45 60 01 f4 00 00 00 00 3e 06 50 9f
0a 01 01 02 1e 01 01 02 5b 9b 30 39 00 00 00 00
00 00 00 00 50 02 00 00 f8 3c 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 ac 7a 04 00
00 00 00 00 b4 e3 15 4e bf 65 06 00 04 22 38 02
38 02 00 00 00 11 00 03 02 00 00 00 50 d0 1a 08
30 de be bf e4 f3 19 08
```

```
user@host> show security datapath-debug counters
Datapath debug counters
Packet Filter 1:
lt-enter
Chassis 0 FPC 0 PIC 1: 0
lt-enter
Chassis 0 FPC 0 PIC 0: 1
lt-leave
Chassis 0 FPC 0 PIC 1: 0
lt-leave
Chassis 0 FPC 0 PIC 0: 1
np-egress
Chassis 0 FPC 1 PIC 3: 0
np-egress
Chassis 0 FPC 1 PIC 1: 0
np-egress
Chassis 0 FPC 1 PIC 2: 0
np-egress
Chassis 0 FPC 1 PIC 0: 1
pot
Chassis 0 FPC 0 PIC 1: 0
pot
Chassis 0 FPC 0 PIC 0: 6
np-ingress
Chassis 0 FPC 1 PIC 3: 0
np-ingress
Chassis 0 FPC 1 PIC 1: 0
```

```
np-ingress
Chassis 0 FPC 1 PIC 2: 0
np-ingress
Chassis 0 FPC 1 PIC 0: 1
lbt
Chassis 0 FPC 0 PIC 1: 0
lbt
Chassis 0 FPC 0 PIC 0: 4
jexec
Chassis 0 FPC 0 PIC 1: 0
jexec
Chassis 0 FPC 0 PIC 0: 4
```

SEE ALSO

Understanding Data Path Debugging for Logical Systems | 520

Troubleshooting DNS Name Resolution in Logical System Security Policies (Primary Administrators Only)

IN THIS SECTION

- Problem | 527
- Cause | 528
- Solution | 528

Problem

Description

The address of a hostname in an address book entry that is used in a security policy might fail to resolve correctly.

Cause

Normally, address book entries that contain dynamic hostnames refresh automatically for SRX Series devices. The TTL field associated with a DNS entry indicates the time after which the entry should be refreshed in the policy cache. Once the TTL value expires, the SRX Series device automatically refreshes the DNS entry for an address book entry.

However, if the SRX Series device is unable to obtain a response from the DNS server (for example, the DNS request or response packet is lost in the network or the DNS server cannot send a response), the address of a hostname in an address book entry might fail to resolve correctly. This can cause traffic to drop as no security policy or session match is found.

Solution

The primary administrator can use the `show security dns-cache` command to display DNS cache information on the SRX Series device. If the DNS cache information needs to be refreshed, the primary administrator can use the `clear security dns-cache` command.

NOTE: These commands are only available to the primary administrator on devices that are configured for logical systems. This command is not available in user logical systems or on devices that are not configured for logical systems.

SEE ALSO

Understanding Logical Systems Security Policies

RELATED DOCUMENTATION

Security Profiles for Logical Systems | 70

3

CHAPTER

Tenant Systems

Tenant Systems Overview | 530

Security Zones for Tenant Systems | 572

Flow for Tenant Systems | 578

Flow Trace for Tenant Systems | 614

Firewall Authentication for Tenant Systems | 617

Security Policies for Tenant Systems | 654

Screen Options for Tenant Systems | 664

NAT for Tenant Systems | 672

UTM for Tenant Systems | 682

IDP for Tenant Systems | 690

ALG for Tenant Systems | 713

DHCP for Tenant Systems | 727

Security Log for Tenant Systems | 737

AppQoS for Tenant Systems | 750

Application Security for Tenant Systems | 758

Tenant Systems Overview

IN THIS SECTION

- Understanding Tenant Systems | 530
- Tenant System Configuration Overview | 537
- Configuring a Routing Instance for a Tenant System | 539
- Understanding Routing and Interfaces for Tenant Systems | 541
- Understanding Tenant System Security Profiles (Primary Administrators Only) | 548
- Example: Creating Tenant Systems, Tenant System Administrators, and an Interconnect VPLS Switch | 554

A tenant system supports routing, services and security features.

Understanding Tenant Systems

IN THIS SECTION

- Differences Between Logical Systems and Tenant Systems | 531
- Use Cases for Logical Systems and Tenant Systems | 532
- Deployment Scenarios for Multitenant Systems | 533
- Benefits of Tenant Systems | 534
- Roles and Responsibilities of Primary Administrator and Tenant System Administrator | 534
- Tenant System Capacity | 536

A tenant system logically partitions the physical firewall into separate and isolated logical firewall. Although similar to logical systems, tenant systems have much higher scalability and fewer routing features. Each tenant system on a device allows you to control a discrete administrative domain for security services. By transforming your device into a multitenant system, you can provide various departments, organizations, customers, and partners—depending on your environment—private and logically separated use of system resources and tenant-specific views of security configuration and KPIs.

A primary administrator creates and manages all the tenant systems. Figure 12 on page 531 shows a single device with a primary logical system and discrete tenant systems.

Figure 12: Tenant Systems

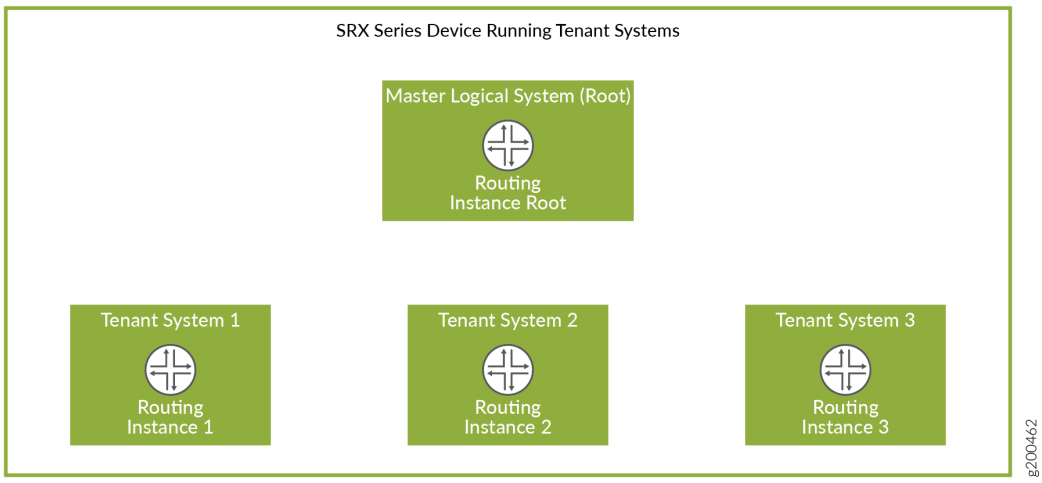

Differences Between Logical Systems and Tenant Systems

Table 27 on page 531 describes the key differences between logical systems and tenant systems.

Table 27: Differences Between Logical Systems and Tenant Systems

Functionality	Logical Systems	Tenant Systems
Feature support	Supports all the routing features to provide optimal data routing paths.	Supports routing features and high-scale security virtualization to isolate customer environments.
Scalability	A maximum of 32 logical systems can be configured on a physical SRX Series device.	A maximum of 500 tenant systems can be configured on a physical SRX Series device to provide high scalability.

Table 27: Differences Between Logical Systems and Tenant Systems *(Continued)*

Functionality	Logical Systems	Tenant Systems
Routing protocol process	Every logical system needs an individual copy of the routing protocol process to logically separate the resources on a device.	The primary logical system has a single routing protocol process, which is shared by the tenant systems. Routing instances supported by this single routing protocol process achieve the security resource separation on the firewall.
Routing instance	A default routing instance is automatically created for every logical system.	Starting in Junos OS Release 19.2R1, the virtual-router configured in a tenant system is passed as the default routing-instance to `ping`, `telnet`, `ssh`, `traceroute`, `show arp`, `clear arp`, `show ipv6 neighbors`, and `clear ipv6 neighbors` commands.
Logical interface configuration	The primary administrator assigns the logical interfaces and the logical system administrator can configure the interface attributes.	A tenant system administrator cannot configure the logical interfaces. The primary administrator assigns the logical interfaces to a tenant system.

Use Cases for Logical Systems and Tenant Systems

A logical system is used when more than one virtual router is required. For example, you have multiple connections to the external network and they cannot co-exist in the same virtual router. Tenant systems are used when you need to separate departments, organization, or customers and each of them can be limited to one virtual router. The main difference between a logical system and a tenant system is that a logical system supports advanced routing functionality using multiple routing instances. In comparison, a tenant system supports only one routing instance, but supports the deployment of significantly more tenants per system.

Deployment Scenarios for Multitenant Systems

You can deploy an SRX Series device running a multitenant system in many environments such as a managed security service provider (MSSP), an enterprise network, or a branch office segment. Table 28 on page 533 describes the various deployment scenarios and the roles played by the tenant systems in such scenarios.

Table 28: Deployment Scenarios with Respect to Tenant Systems

Deployment Scenarios	Roles of a Tenant System
Managed security service provider (MSSP)	• In a managed security service provider (MSSP), each customer can be isolated from other customer to protect data privacy. Customers that require defined service level agreements (SLAs) can be be allocated memory and system resources to meet these SLAs. • The customer can configure distinct security policies for compliance and control per tenant system.
Enterprise network	• A tenant system can be assigned to a workgroup, department, or other organizational construct within an enterprise. • A tenant system can define the distinct security policies for the enterprise workgroup, department, or other organizational construct of the enterprise.
Branch office segment	• In a branch office, a tenant system can individually manage and segregate corporate and guest traffic. • Advanced security policies can be configured per tenant system; this approach allows granular control of the security policies. • A tenant system provides ease of management and troubleshooting.

Benefits of Tenant Systems

- Curtail cost by reducing the number of physical devices required for your organization. You can consolidate services for various groups of users on a single device and reduce the hardware costs, power expenditure, and rack space.
- Provide isolation and logical separation at the tenant system level. Provides the ability to separate tenant systems with administrative separation at large scale in which each tenant system can define its own security controls and restrictions without impacting other tenant systems.

Roles and Responsibilities of Primary Administrator and Tenant System Administrator

A primary administrator creates and manages all the tenant systems. A primary logical system is created at the root level and is allocated a single routing protocol process. Although this routing protocol process is shared, tenant systems enable logical resource separation on the firewall. By default, all system resources are assigned to the primary logical system, and the primary administrator allocates them to the tenant system administrators.

NOTE: In Junos OS command-line reference, primary logical system is referred as root logical system.

A tenant system is created that is subtended by the primary logical system. Although all the tenants under the primary logical system share a single routing process, each tenant system has a single routing instance. Table 29 on page 535 describes the roles and responsibilities of the primary administrator and tenant system administrator.

Table 29: Roles and Responsibilities With Respect to Tenant Systems

Roles	Definition	Responsibilities
Primary administrator	A user account with superuser configuration and verification privileges for all logical systems and tenant systems.	• View and access all logical systems and tenant systems. • Create login accounts for all the tenant systems and assign the login accounts to the appropriate tenant system. • Create and allocate the resources to the tenant systems. • Create one custom routing instance under the tenant system which acts as the default routing instance for the tenant system. • Create a virtual router under the tenant system and assign it to the tenant system. • Create logical interfaces to assign to the tenant systems. • Manage the tenant systems in the primary logical system. • Ensure duplicate names for tenant system, logs, and trace file do not exist.

Table 29: Roles and Responsibilities With Respect to Tenant Systems *(Continued)*

Roles	Definition	Responsibilities
Tenant system administrator	A tenant system account with all configuration and verification privileges. **NOTE**: The configuration and verification privileges of a tenant system administrator depends on the permission assigned to them by the primary administrator while creating the tenant system administrator. Multiple tenant system administrators can be created for a tenant system with different permission levels based on your requirement.	• Access and view the resources of the tenant system. • Configure the resources allocated and routing protocols. • Configure schedulers, security profiles, and security features. The following privileges are not supported by the tenant system administrator: • Define access restrictions and the default routing instance for the tenant system. • Access and view the resources of other tenant systems. • Modify the number of allocated resources for a tenant system. • Create logical interfaces, virtual router, and policy options.

Tenant System Capacity

The maximum number of tenant systems that can be created on the device are listed in Table 30 on page 536.

Table 30: Tenant Systems Capacity

Platform	Logical Systems Capacity	Tenant Systems Capacity for Junos OS Release 18.4R1
SRX1500	32	50
SRX4100 and SRX4200	32	200
SRX4600	32	300

Table 30: Tenant Systems Capacity *(Continued)*

Platform	Logical Systems Capacity	Tenant Systems Capacity for Junos OS Release 18.4R1
SRX5400, SRX5600, and SRX5800 Series devices with SPC2 cards	32	100
SRX5400, SRX5600, and SRX5800 Series devices with SPC3 cards	32	500
SRX5400, SRX5600, and SRX5800 Series devices with SPC2 and SPC3 cards	32	100

Starting in Junos OS Release 18.4R1, tenant systems can be supported on an SRX5000 series security services gateway equipped with a combination of third generation service processing cards (SRX5K-SPC3) and second generation service processing cards (SRX5K-SPC-4-15-320). Prior to Junos OS Release 18.4R1, tenant systems was supported on SPC2 only.

SEE ALSO

Primary Logical Systems Overview | 23

Tenant System Configuration Overview

The primary administrator creates a tenant system and assigns an administrator for managing the tenant system. A tenant system can have multiple administrators. The roles and responsibilities of a tenant system administrator are explained in "Understanding Tenant Systems" on page 530.

The primary administrator configures the logical interfaces and assigns those interfaces to the tenant system. Configure one routing instance and the routing protocols, and add options for the routing instance. See "Configuring a Routing Instance for a Tenant System" on page 539.

Tenant systems have their own configuration database. After successful configuration, the changes are merged to the primary database for each tenant systems. Multiple tenant systems can perform configuration changes at a time. You can commit the changes for only one tenant at a time. If the primary administrator and a tenant system administrator performs configuration changes simultaneously, the configuration changes performed by the primary administrator override the configuration changes performed by the tenant system administrator.

The following steps explain the tasks that the tenant system administrator performs to configure the security features in a tenant system:

1. Use the SSH service to access the device, and then log in to the tenant system with the login ID and password provided by the primary administrator.

```
login: <tenant_name>
password: <password>
```

   After you are authenticated, the presence of the “>” prompt indicates that you accessed to the CLI operational mode. The prompt is preceded by a string that contains the username, the hostname of the device, and the name of the tenant system. When the CLI starts, you are at the top level in operational mode.

```
TSYS1_admin1@host:TSYS1>
```

2. Access the configuration mode by entering the `configure` command.

```
TSYS1_admin1@host:TSYS1> configure
TSYS1_admin1@host:TSYS1#
```

3. Enter the **quit** command to exit the configuration mode and return to the CLI operational mode.

```
TSYS1_admin1@host:TSYS1# quit
TSYS1_admin1@host:TSYS1>
```

4. Configure the following security features in the tenant system as necessary:

   - Create zones for the tenant system and bind the logical interfaces to the zones. Create address books and use them in the security policies. See "Example: Configuring Zones in the Tenant System" on page 574.

   - Configure screen options at the zone level. See Example: Configuring Screen Options for a Tenant System.

   - Configure security policies between zones in the tenant system. See Example: Configuring Security Policies in a Tenant System.

     Custom applications or application sets can be created for specific types of traffic. To create a custom application, use the `application` configuration statement at the [`edit applications`] hierarchy level. To create an application set, use the `application-set` configuration statement at the [`edit applications`] hierarchy level.

- Configure firewall authentication to the tenant system. The primary administrator creates access profiles in the primary logical system. The tenant system administrator then configures a security policy that specifies firewall authentication for matching traffic and configures the type of authentication (pass-through or Web authentication), default access profile, and success banner. See "Configuring Firewall Authentication for a Tenant System" on page 620.
- Configure Network Address Translation (NAT) for the tenant system. See Example: Configuring Network Address Translation for the Tenant Systems.
- Configure Application Layer Gateway (ALG) for the tenant system. See Example: Configuring ALG in Tenant System.
- Configure Intrusion Detection and Prevention (IDP) policies and attacks for the tenant system. See "Example: Configuring IDP Policies and Attacks for Tenant Systems" on page 694.

Configuring a Routing Instance for a Tenant System

A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. A set of interfaces that belong to the routing instance and the routing protocol parameters control the information in the routing instance. A tenant system can configure the assigned routing instance and the interfaces that belong to the routing instance within a tenant system.

NOTE: Only one routing instance can be created for a tenant system.

The following procedure describes the steps to configure a routing instance and interfaces in a routing table for a tenant system:

1. Create a tenant system named `TSYS1`.

```
[edit]
user@host# set tenants TSYS1
```

2. Create a routing instance `r1` and assign the routing instance type for the tenant system.

```
[edit]
user@host# set tenants TSYS1 routing-instances r1 instance-type virtual-router
```

3. Specify the interface name for the routing instance.

```
[edit]
user@host# set tenants TSYS1 routing-instances r1 interface lt-0/0/0.101
user@host# set tenants TSYS1 routing-instances r1 interface xe-0/0/0.0
user@host# set tenants TSYS1 routing-instances r1 interface xe-0/0/1.0
```

4. Specify the routing option for the routing instance.

```
[edit]
user@host# set tenants TSYS1 routing-instances r1 routing-options router-id 1.1.1.101
```

5. Commit the configuration.

```
[edit]
user@host# commit
```

To view the configuration for the tenant system `TSYS1`, run the `show tenants TSYS1` command.

```
routing-instances {
    r1 {
        instance-type virtual-router;
        interface lt-0/0/0.101;
        interface xe-0/0/0.0;
        interface xe-0/0/1.0;
        routing-options {
            router-id 1.1.1.101;
        }
    }
}
```

The `show tenants TSYS1` command displays all the routing instance parameters configured for the tenant system `TSYS1`.

Understanding Routing and Interfaces for Tenant Systems

IN THIS SECTION

- Overview: Configuring Routing and Interfaces for Tenant Systems | 541

A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. The interfaces are used for forwarding data for the routing instance, and to learn the routing information from other peers (SRX devices) using routing protocols.

A Logical interface (IFL) can be defined at either one of the following levels:

- Global level (root logical system)
- User logical system level
- Tenant system level (Starting from Release Junos OS 18.4R1)

The IFL defined at the global level can be used either in root logical system or in one of the tenant systems. The IFL defined in a tenant system can be used in that tenant system only.

Default routing instance is not available for tenant systems. So, when a custom routing instance is created for a tenant system, all the interfaces defined in that tenant system should be added to that routing instance.

Overview: Configuring Routing and Interfaces for Tenant Systems

IN THIS SECTION

- Requirements | 542
- Overview | 542
- Configuration | 542

This overview shows how to configure interfaces and routing instances for a tenant system.

Requirements

Before you begin:

- Determine which logical interfaces and, optionally, which logical tunnel interfaces are allocated. See Tenant System Configuration Overview.

Overview

The following procedure describes the steps to configure a routing instance and interfaces in a routing table within a tenant system.

This topic configures the interfaces and routing instances described in Table 31 on page 542.

Table 31: User Tenant System Interface and Routing Instance Configuration

Feature	Name	Configuration Parameters
Interface	ge-0/0/2.1 ge-0/0/2.2 ge-0/0/2.3	• IP address 10.0.0.1/24 • IP address 10.0.0.2/24 • IP address 10.0.0.3/24
Routing instance	r1 r2	• Instance type: virtual router • Includes interfaces ge-0/0/2.1, ge-0/0/2.3, and ge-0/0/2.2

Configuration

IN THIS SECTION

- Procedure | 543

Procedure

CLI Quick Configuration

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

```
set interfaces ge-0/0/2 vlan-tagging
set interfaces ge-0/0/2.3 vlan-id 103
set interfaces ge-0/0/2.3 family inet address 10.0.0.3/24
set tenants TSYS1
set tenants TSYS1 interfaces ge-0/0/2.1 vlan-id 101
set tenants TSYS1 interfaces ge-0/0/2.1 family inet address 10.0.0.1/24
set tenants TSYS1 routing-instances r1 instance-type virtual-router
set tenants TSYS1 routing-instances r1 interface ge-0/0/2.1
set tenants TSYS1 routing-instances r1 interface ge-0/0/2.3
set tenants TSYS2
set tenants TSYS2 interfaces ge-0/0/2.2 vlan-id 102
set tenants TSYS2 interfaces ge-0/0/2.2 family inet address 10.0.0.2/24
set tenants TSYS2 routing-instances r2 instance-type virtual-router
set tenants TSYS2 routing-instances r2 interface ge-0/0/2.2
```

Step-by-Step Procedure

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

To configure an interface and a routing instance in a user logical system:

1. Configure the interfaces to support VLAN tagging.

```
[edit]
user@host# set interfaces ge-0/0/2 vlan-tagging
```

2. Configure the IFL at the root level.

```
[edit]
set interfaces ge-0/0/2.3 vlan-id 103
set interfaces ge-0/0/2.3 family inet address 10.0.0.3/24
```

3. Create a tenant system named TSYS1.

```
[edit]
user@host# set tenants TSYS1
```

4. Define the Interface in the tenant system TSYS1.

```
[edit]
user@host# set tenants TSYS1 interfaces ge-0/0/2.1 vlan-id 101
user@host# set tenants TSYS1 interfaces ge-0/0/2.1 family inet address 10.0.0.1/24
user@host# set tenants TSYS1 routing-instances r1 interface ge-0/0/2.3
```

5. Create a routing instance r1 and assign the routing instance type for the tenant system.

```
[edit]
user@host# set tenants TSYS1 routing-instances r1 instance-type virtual-router
```

6. Specify the interface name for the routing instance.

```
[edit]
user@host# set tenants TSYS1 routing-instances r1 interface ge-0/0/2.1
```

7. Create a tenant system named TSYS2.

```
[edit]
user@host# set tenants TSYS2
```

8. Define the Interface in the tenant system TSYS2.

```
[edit]
user@host# set tenants TSYS2 interfaces ge-0/0/2.2 vlan-id 102
user@host# set tenants TSYS2 interfaces ge-0/0/2.2 family inet address 10.0.0.2/24
```

9. Create a routing instance `r2` and assign the routing instance type for the tenant system.

```
[edit]
user@host# set tenants TSYS2 routing-instances r2 instance-type virtual-router
```

10. Specify the interface name for the routing instance.

```
[edit]
user@host# set tenants TSYS2 routing-instances r2 interface ge-0/0/2.2
```

11. Commit the configuration.

```
[edit]
user@host# commit
```

Results

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

```
[edit]
user@host# show interfaces
ge-0/0/2 {
    vlan-tagging;
    unit 3 {
        vlan-id 103;
        family inet {
            address 10.0.0.3/24;
        }
```

```
    }
}
```

```
[edit]
user@host# show tenants
TSYS1 {
    interfaces {
        ge-0/0/2 {
            unit 1 {
                vlan-id 101;
                family inet {
                    address 10.0.0.1/24;
                }
            }
        }
    }
    routing-instances {
        r1 {
            instance-type virtual-router;
            interface ge-0/0/2.1;
            interface ge-0/0/2.3;
        }
    }
}
TSYS2 {
    interfaces {
        ge-0/0/2 {
            unit 2 {
                vlan-id 102;
                family inet {
                    address 10.0.0.2/24;
                }
            }
        }
    }
    routing-instances {
        r2 {
            instance-type virtual-router;
            interface ge-0/0/2.2;
        }
```

```
    }
}
```

The `show tenants` command displays all the interfaces that are defined in the tenant systems `TSYS1` and `TSYS2`, and the routing instance parameters configured for both the tenant systems.

```
user@host> show interfaces ge-0/0/2.1 detail
  Logical interface ge-0/0/2.1 (Index 89) (SNMP ifIndex 548) (Generation 161)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.101 ]  Encapsulation: ENET2
    Tenant Name: TSYS1
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :                  46
     Input  packets:                   0
     Output packets:                   1
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                  46
     Input  packets:                   0
     Output packets:                   1
    Transit statistics:
     Input  bytes  :                   0                  0 bps
     Output bytes  :                   0                  0 bps
     Input  packets:                   0                  0 pps
     Output packets:                   0                  0 pps
    Security: Zone: Null
    Flow Statistics :
.............................
```

```
user@host> show interfaces ge-0/0/2.2 detail
  Logical interface ge-0/0/2.2 (Index 90) (SNMP ifIndex 549) (Generation 162)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.102 ]  Encapsulation: ENET2
    Tenant Name: TSYS2
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :                  46
     Input  packets:                   0
     Output packets:                   1
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                  46
```

```
 Input  packets:                  0
 Output packets:                  1
Transit statistics:
 Input  bytes  :                  0                   0 bps
 Output bytes  :                  0                   0 bps
 Input  packets:                  0                   0 pps
 Output packets:                  0                   0 pps
Security: Zone: Null
Flow Statistics :
Flow Input statistics :
  Self packets :                   0
  ICMP packets :                   0
  VPN packets :              ..............................
```

SEE ALSO

Tenant System Configuration Overview

Understanding Tenant System Security Profiles (Primary Administrators Only)

IN THIS SECTION

- Tenant Systems Security Profiles | **549**
- Understanding How the System Assesses Resources Assignment and Use Across the Tenant Systems | **549**
- Cases: Assessments of Reserved Resources Assigned Through Security Profiles | **551**

Tenant systems allow you to virtually divide a supported SRX Series device into multiple devices, securing them from intrusion and attacks, and protecting them from faulty conditions outside their own contexts. To protect tenant systems, security resources are configured in a manner similar to how they are configured for a discrete device. However, the primary administrator assigns resources to the tenant systems.

An SRX Series device running tenant systems can be partitioned into tenant systems, an interconnected tenant system, if necessary, and the default primary logical system. When the system is initialized, the primary logical system is created at the root. All system resources are assigned to it, effectively creating a default primary logical system security profile. To distribute security resources across the tenant systems, the primary administrator creates security profiles that specify the resources to be allocated to a tenant system. Only the primary administrator can configure security profiles and bind them to the tenant systems. The tenant system administrator uses these resources for the respective tenant system.

The tenant systems are defined by the resources allocated to them, including security components, interfaces, routing instance, static routes, and dynamic routing protocols. The primary administrator configures the security profiles and assigns them to the tenant systems. You cannot commit a tenant system configuration without a security profile assigned to it.

This topic includes the following sections:

Tenant Systems Security Profiles

The primary administrator can configure and assign a security profile to a specific tenant system or multiple tenant systems. The maximum number of security profiles that can be configured depends on the capacity of an SRX Series device. When the maximum number of security profiles have been created, you need to delete a security profile and commit the configuration change before you can create and commit another security profile. In many cases, fewer security profiles are needed because you can bind a single security profile to more than one tenant system.

Security profiles allow you to:

- Share the device's resources, including policies, zones, addresses and address books, flow sessions, and various forms of NAT, among all tenant systems appropriately. You can assign various amounts of a resource to the tenant systems and allow the tenant systems to utilize the resources effectively.

  Security profiles protect against one tenant system exhausting a resource that is required at the same time by other tenant systems. Security profiles protect critical system resources and maintain a better performance among tenant systems when the device is experiencing a heavy traffic flow. Security profiles defend against one tenant system dominating the use of resources and allow the other tenant systems to use the resources effectively.

- Configure the device in a scalable way to allow for creation of additional tenant systems.

You need to delete the security profile of a tenant system before you can delete the tenant system.

Understanding How the System Assesses Resources Assignment and Use Across the Tenant Systems

To provision a tenant system with security features, the primary administrator configures a security profile that specifies the resource for each security feature:

- A reserved quota that guarantees that the specified resource amount is always available to the tenant system.
- A maximum allowed quota. If a tenant system requires additional resources that exceed the reserved quota, then it can utilize the resources configured for the global maximum amount if the global resources are not allocated to the other tenant systems. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. The tenant systems need to utilize the global resources effectively based on the available resources.

If a reserved quota is not configured for a resource, the default value is 0. If a maximum allowed quota is not configured for a resource, the default value is the global system quota for the resource (global system quotas are platform-dependent). The primary administrator must configure the appropriate maximum allowed quota values in the security profiles so that the maximum resource usage of a specific tenant system does not negatively impact other tenant systems configured on the device.

The system maintains a count of all allocated resources that are reserved, used, and made available again when a tenant system is deleted. This count determines whether resources are available to use for tenant systems or to increase the amount of the resources allocated to existing tenant systems through their security profiles.

Resources configured in security profiles are characterized as static modular resources or dynamic resources. For static resources, we recommend setting a maximum quota for a resource equal or close to the amount specified as its reserved quota, to allow for scalable configuration of tenant systems. A maximum quota for a resource gives a tenant system greater flexibility through access to a larger amount of that resource, but it constrains the amount of resources available to allocate to other tenant systems.

The following security features resources can be specified in a security profile:

- Security zones
- Addresses and address books for security policies
- Application firewall rule sets
- Application firewall rules
- Firewall authentication
- Flow sessions and gates
- NAT, including:
    - Cone NAT bindings
    - NAT destination rule
    - NAT destination pool

- NAT IP address in source pool without Port Address Translation (PAT)

  NOTE: IPv6 addresses in IPv6 source pools without PAT are not included in security profiles.

- NAT IP address in source pool with PAT
- NAT port overloading
- NAT source pool
- NAT source rule
- NAT static rule

NOTE: All resources except flow sessions are static.

You can modify a tenant system security profile dynamically while the security profile is assigned to other tenant systems. However, to ensure that the system resource quota is not exceeded, the system takes the following actions:

- If a static quota is changed, the system process that maintains the tenant system counts for resources specified in security profiles subsequently reevaluates the security profiles assigned to the profile associated with the static quota. This check identifies the number of resources assigned across all tenant systems to determine whether the allocated resources, including their increased amounts are available.

  These quota checks are the same quota checks that the system performs when you add a tenant system and bind a security profile to it. They are also performed when you bind a different security profile from the security profile that is currently assigned to it to an existing tenant system (or the primary logical system).

- If a dynamic quota is revised, no check is performed, but the revised quota is imposed on future resource usage.

Cases: Assessments of Reserved Resources Assigned Through Security Profiles

To understand how the system assesses allocation of reserved resources through security profiles, consider the following three cases explained in Table 33 on page 553 and that address allocation of the resources and zones. To keep the example simple, 10 zones are allocated in security-profile-1: 4 reserved zones and 6 maximum zones. This example assumes that the maximum amount specified—six zones—is available for the tenant systems. The system maximum number of zones is 10.

The three cases address the configuration across the tenant systems. The three cases verify whether a configuration succeeds or fails when it is committed based on the allocation of zones.

Table 32 on page 552 shows the security profiles and their zone allocations.

Table 32: Security Profiles Used for Reserved Resource Assessments

Two Security Profiles Used in the Configuration Cases
security-profile-1 • zones reserved quota = 4 • zones maximum quota = 6 **NOTE**: The primary administrator dynamically increases the reserved zone count specified in this profile later.
primary-logical-system-profile • zones maximum quota = 10 • no reserved quota

Table 33 on page 553 shows three cases that illustrate how the system assesses reserved resources for zones across the tenant systems based on the security profile configurations.

- The configuration for the first case succeeds because the cumulative reserved resource quota for zones configured in the security profiles bound to all tenant systems is 8, which is less than the system maximum resource quota.
- The configuration for the second case fails because the cumulative reserved resource quota for zones configured in the security profiles bound to all logical systems is 12, which is greater than the system maximum resource quota.
- The configuration for the third case fails because the cumulative reserved resource quota for zones configured in the security profiles bound to all tenant systems is 12, which is greater than the system maximum resource quota.

Table 33: Reserved Resource Allocation Assessment Across Tenant Systems

Reserved Resource Quota Checks Across Tenant Systems
Example 1: Succeeds This configuration is within bounds: 4+4+0=8, maximum capacity =10. Security Profiles Used • The security profile security-profile-1 is bound to two tenant systems: tenant-system-1 and tenant-system-2. • The primary-logical-system-profile profile is used exclusively for the primary logical system. • tenant-system-1 = 4 reserved zones. • tenant-system-2 = 4 reserved zones. • primary-logical-system = 0 reserved zones.
Example 2: Fails This configuration is out of bounds: 4+4+4=12, maximum capacity =10. • tenant-system-1 = 4 reserved zones. • tenant-system-2 = 4 reserved zones. • primary-logical-system = 0 reserved zones. • new-tenant-system = 4 reserved zones. Security Profiles • The security profile security-profile-1 is bound to two tenant systems: tenant-system-1 and tenant-system-2. • The primary-logical-system-profile is bound to the primary logical system and used exclusively for it. • The primary administrator configures a new tenant system called new-tenant-system and binds security-profile-1 to it.

Table 33: Reserved Resource Allocation Assessment Across Tenant Systems *(Continued)*

Reserved Resource Quota Checks Across Tenant Systems
Example 3: Fails This configuration is out of bounds: 6+6=12, maximum capacity =10. The primary administrator modifies the reserved zones quota in security-profile-1, increasing the count to 6. • tenant-system-1 = 6 reserved zones. • tenant-system-2 = 6 reserved zones. • primary-logical-system = 0 reserved zones.

Example: Creating Tenant Systems, Tenant System Administrators, and an Interconnect VPLS Switch

IN THIS SECTION

- Requirements | **554**
- Overview | **555**
- Full SRX Quick Configuration | **556**
- Verification | **569**

This example shows how to create tenant systems, tenant system administrators, and an interconnect VPLS switch. Only the primary administrator can create user login accounts for tenant system administrators and interconnect VPLS switch.

Requirements

This example uses the following hardware and software components:

- SRX Series device.
- Junos OS Release 18.4R1 and later releases.

- "VSRX requires 20.1R1 and 16GB of memory." on page 21

- Before you begin creating the tenant systems, tenant system administrators, and an interconnect VPLS switch, read "Tenant Systems Overview" on page 530 to understand how this task fits into the overall configuration process.

Overview

IN THIS SECTION

- Topology | 555

This example shows how to create the tenant systems TSYS1, TSYS2, and TSYS3, and the tenant system administrators for them. You can create multiple tenant system administrators for a tenant system with different permission levels based on your requirements.

This topic also covers the interconnect virtual private LAN service (VPLS) switch connecting one tenant system to another on the same device. The VPLS switch enables both transit traffic and traffic terminated at a tenant system to pass between tenant systems. To allow traffic to pass between tenant systems, logical tunnel (lt-0/0/0) interfaces should be configured in the same subnet.

Topology

The Figure 13 on page 556 shows an SRX Series device deployed and configured for tenant systems. The configuration example uses static routing to allow the PCs to reach the Internet.

Figure 13: Creating Tenant Systems and Interconnect VPLS Switch

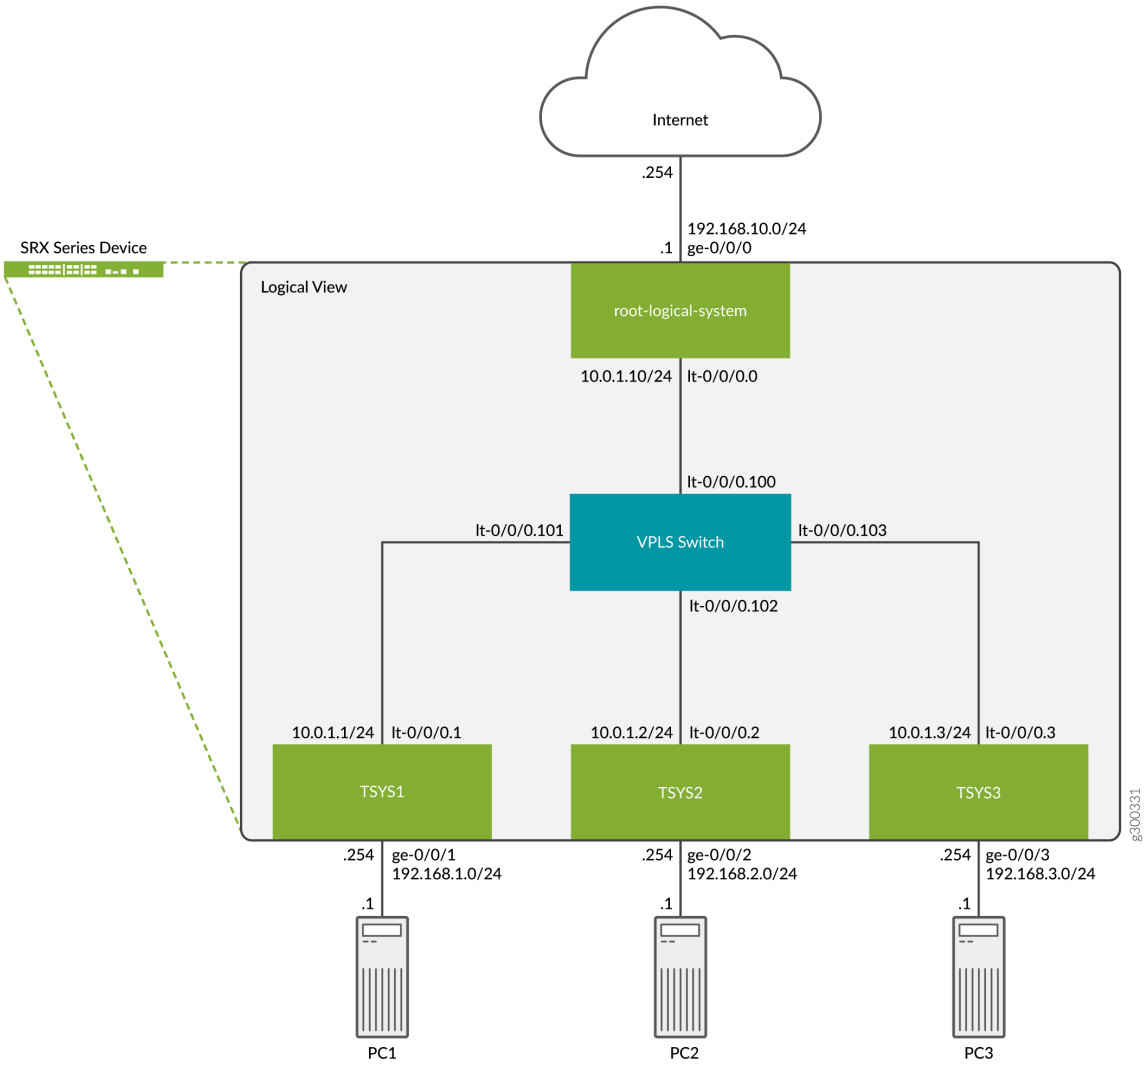

Full SRX Quick Configuration

IN THIS SECTION

- Configuring Logical and Tenant Systems, and Interconnect VPLS Switch | 557

CLI Quick Configuration

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

```
set system login class TSYS1admin1 tenant TSYS1
set system login class TSYS1admin1 permissions all
set system login class TSYS2admin1 tenant TSYS2
set system login class TSYS2admin1 permissions all
set system login class TSYS3admin1 tenant TSYS3
set system login class TSYS3admin1 permissions all
set system login user TSYS1admin1 uid 2001
set system login user TSYS1admin1 class TSYS1admin1
set system login user TSYS1admin1 authentication encrypted-password "$ABC123"
set system login user TSYS2admin1 uid 2003
set system login user TSYS2admin1 class TSYS2admin1
set system login user TSYS2admin1 authentication encrypted-password "$ABC123"
set system login user TSYS3admin1 uid 2005
set system login user TSYS3admin1 class TSYS3admin1
set system login user TSYS3admin1 authentication encrypted-password "$ABC123"
set system security-profile SP0 logical-system root-ls
set system security-profile SP1 tenant TSYS1
set system security-profile SP2 tenant TSYS2
set system security-profile SP3 tenant TSYS3
set logical-systems root-ls interfaces ge-0/0/0 unit 0 family inet address 192.168.10.1/24
set logical-systems root-ls interfaces lt-0/0/0 unit 0 encapsulation ethernet
set logical-systems root-ls interfaces lt-0/0/0 unit 0 peer-unit 100
set logical-systems root-ls interfaces lt-0/0/0 unit 0 family inet address 10.0.1.10/24
set logical-systems root-ls routing-options static route 192.168.1.0/24 next-hop 10.0.1.1
set logical-systems root-ls routing-options static route 192.168.2.0/24 next-hop 10.0.1.2
set logical-systems root-ls routing-options static route 192.168.3.0/24 next-hop 10.0.1.3
set logical-systems root-ls security address-book global address TSYS1 192.168.1.0/24
set logical-systems root-ls security address-book global address TSYS2 192.168.2.0/24
set logical-systems root-ls security address-book global address TSYS3 192.168.3.0/24
set logical-systems root-ls security policies from-zone trust to-zone untrust policy allow-out
match source-address TSYS1
set logical-systems root-ls security policies from-zone trust to-zone untrust policy allow-out
match source-address TSYS2
```

```
set logical-systems root-ls security policies from-zone trust to-zone untrust policy allow-out
match source-address TSYS3
set logical-systems root-ls security policies from-zone trust to-zone untrust policy allow-out
match destination-address any
set logical-systems root-ls security policies from-zone trust to-zone untrust policy allow-out
match application any
set logical-systems root-ls security policies from-zone trust to-zone untrust policy allow-out
then permit
set logical-systems root-ls security zones security-zone trust host-inbound-traffic system-
services ping
set logical-systems root-ls security zones security-zone trust interfaces lt-0/0/0.0
set logical-systems root-ls security zones security-zone untrust host-inbound-traffic system-
services ping
set logical-systems root-ls security zones security-zone untrust interfaces ge-0/0/0.0
set interfaces lt-0/0/0 unit 1 encapsulation ethernet
set interfaces lt-0/0/0 unit 1 peer-unit 101
set interfaces lt-0/0/0 unit 1 family inet address 10.0.1.1/24
set interfaces lt-0/0/0 unit 2 encapsulation ethernet
set interfaces lt-0/0/0 unit 2 peer-unit 102
set interfaces lt-0/0/0 unit 2 family inet address 10.0.1.2/24
set interfaces lt-0/0/0 unit 3 encapsulation ethernet
set interfaces lt-0/0/0 unit 3 peer-unit 103
set interfaces lt-0/0/0 unit 3 family inet address 10.0.1.3/24
set interfaces lt-0/0/0 unit 100 encapsulation ethernet-vpls
set interfaces lt-0/0/0 unit 100 peer-unit 0
set interfaces lt-0/0/0 unit 101 encapsulation ethernet-vpls
set interfaces lt-0/0/0 unit 101 peer-unit 1
set interfaces lt-0/0/0 unit 102 encapsulation ethernet-vpls
set interfaces lt-0/0/0 unit 102 peer-unit 2
set interfaces lt-0/0/0 unit 103 encapsulation ethernet-vpls
set interfaces lt-0/0/0 unit 103 peer-unit 3
set interfaces ge-0/0/1 unit 0 family inet address 192.168.1.254/24
set interfaces ge-0/0/2 unit 0 family inet address 192.168.2.254/24
set interfaces ge-0/0/3 unit 0 family inet address 192.168.3.254/24
set routing-instances VPLS instance-type vpls
set routing-instances VPLS interface lt-0/0/0.100
set routing-instances VPLS interface lt-0/0/0.101
set routing-instances VPLS interface lt-0/0/0.102
set routing-instances VPLS interface lt-0/0/0.103
set tenants TSYS1 routing-instances vr1 instance-type virtual-router
set tenants TSYS1 routing-instances vr1 routing-options static route 0.0.0.0/0 next-hop 10.0.1.10
set tenants TSYS1 routing-instances vr1 interface lt-0/0/0.1
set tenants TSYS1 routing-instances vr1 interface ge-0/0/1.0
```

```
set tenants TSYS1 security address-book global address PC1 192.168.1.0/24
set tenants TSYS1 security policies from-zone PC1 to-zone VPLS policy allow-out match source-
address PC1
set tenants TSYS1 security policies from-zone PC1 to-zone VPLS policy allow-out match
destination-address any
set tenants TSYS1 security policies from-zone PC1 to-zone VPLS policy allow-out match
application any
set tenants TSYS1 security policies from-zone PC1 to-zone VPLS policy allow-out then permit
set tenants TSYS1 security zones security-zone PC1 host-inbound-traffic system-services ping
set tenants TSYS1 security zones security-zone PC1 interfaces ge-0/0/1.0
set tenants TSYS1 security zones security-zone VPLS host-inbound-traffic system-services ping
set tenants TSYS1 security zones security-zone VPLS interfaces lt-0/0/0.1
set tenants TSYS2 routing-instances vr2 instance-type virtual-router
set tenants TSYS2 routing-instances vr2 routing-options static route 0.0.0.0/0 next-hop 10.0.1.10
set tenants TSYS2 routing-instances vr2 interface lt-0/0/0.2
set tenants TSYS2 routing-instances vr2 interface ge-0/0/2.0
set tenants TSYS2 security address-book global address PC2 192.168.2.0/24
set tenants TSYS2 security policies from-zone PC2 to-zone VPLS policy allow-out match source-
address PC2
set tenants TSYS2 security policies from-zone PC2 to-zone VPLS policy allow-out match
destination-address any
set tenants TSYS2 security policies from-zone PC2 to-zone VPLS policy allow-out match
application any
set tenants TSYS2 security policies from-zone PC2 to-zone VPLS policy allow-out then permit
set tenants TSYS2 security zones security-zone PC2 host-inbound-traffic system-services ping
set tenants TSYS2 security zones security-zone PC2 interfaces ge-0/0/2.0
set tenants TSYS2 security zones security-zone VPLS host-inbound-traffic system-services ping
set tenants TSYS2 security zones security-zone VPLS interfaces lt-0/0/0.2
set tenants TSYS3 routing-instances vr3 instance-type virtual-router
set tenants TSYS3 routing-instances vr3 routing-options static route 0.0.0.0/0 next-hop 10.0.1.10
set tenants TSYS3 routing-instances vr3 interface lt-0/0/0.3
set tenants TSYS3 routing-instances vr3 interface ge-0/0/3.0
set tenants TSYS3 security address-book global address PC3 192.168.3.0/24
set tenants TSYS3 security policies from-zone PC3 to-zone VPLS policy allow-out match source-
address PC3
set tenants TSYS3 security policies from-zone PC3 to-zone VPLS policy allow-out match
destination-address any
set tenants TSYS3 security policies from-zone PC3 to-zone VPLS policy allow-out match
application any
set tenants TSYS3 security policies from-zone PC3 to-zone VPLS policy allow-out then permit
set tenants TSYS3 security zones security-zone PC3 host-inbound-traffic system-services ping
set tenants TSYS3 security zones security-zone PC3 interfaces ge-0/0/3.0
```

```
set tenants TSYS3 security zones security-zone VPLS host-inbound-traffic system-services ping
set tenants TSYS3 security zones security-zone VPLS interfaces lt-0/0/0.3
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the CLI User Guide. We will only be covering the configuration of one tenant for the step-by-step procedure.

1. Create the login user accounts for each tenant. We will only show the steps for creating the tenant `TSYS1` user account.

   a. Create the user login class and assign it to the tenant system.

   ```
   [edit]
   user@SRX# set system login class TSYS1admin1 tenant TSYS1
   ```

   b. Assign a permissions level to the login class, for this example we will use the level `all` which allows full access to the tenant system administrator.

   ```
   [edit]
   user@SRX# set system login class TSYS1admin1 permissions all
   ```

   c. Create a user account and assign it to the class from the previous steps. This will allow the user to login to the tenant system.

   ```
   [edit]
   user@SRX# set system login user TSYS1admin1 class TSYS1admin1
   ```

   d. Create a user login password for the user account.

   ```
   [edit]
   user@SRX# set system login user TSYS1admin1 authentication plain-text-password
   New password: "$ABC123"
   Retype new password: "$ABC123"
   ```

2. Configure the VPLS switch. The VPLS switch enables both transit traffic and traffic terminated at a tenant system to pass between tenant systems with a single logical tunnel. Logical tunnel interfaces should be configured in the same subnet to allow traffic between tenant systems.

a. Configure the logical tunnel interfaces.

```
[edit]
user@SRX# set interfaces lt-0/0/0 unit 100 encapsulation ethernet-vpls
user@SRX# set interfaces lt-0/0/0 unit 100 peer-unit 0
user@SRX# set interfaces lt-0/0/0 unit 101 encapsulation ethernet-vpls
user@SRX# set interfaces lt-0/0/0 unit 101 peer-unit 1
user@SRX# set interfaces lt-0/0/0 unit 102 encapsulation ethernet-vpls
user@SRX# set interfaces lt-0/0/0 unit 102 peer-unit 2
user@SRX# set interfaces lt-0/0/0 unit 103 encapsulation ethernet-vpls
user@SRX# set interfaces lt-0/0/0 unit 103 peer-unit 3
```

b. Configure a routing instance for the VPLS switch and assign the logical tunnel interfaces.

```
[edit]
user@SRX# set routing-instances VPLS instance-type vpls
user@SRX# set routing-instances VPLS interface lt-0/0/0.100
user@SRX# set routing-instances VPLS interface lt-0/0/0.101
user@SRX# set routing-instances VPLS interface lt-0/0/0.102
user@SRX# set routing-instances VPLS interface lt-0/0/0.103
```

3. Configure the tenant systems. We are only showing the configuration for one tenant.

a. Configure the interfaces associated with the tenant.

```
[edit]
user@SRX# set interfaces lt-0/0/0 unit 1 encapsulation ethernet
user@SRX# set interfaces lt-0/0/0 unit 1 peer-unit 101
user@SRX# set interfaces lt-0/0/0 unit 1 family inet address 10.0.1.1/24
user@SRX# set interfaces ge-0/0/1 unit 0 family inet address 192.168.1.254/24
```

b. Configure the tenant, routing instance, static routing and assign the interfaces.

```
[edit]
user@SRX# set tenants TSYS1 routing-instances vr1 instance-type virtual-router
user@SRX# set tenants TSYS1 routing-instances vr1 routing-options static route 0.0.0.0/0
next-hop 10.0.1.10
user@SRX# set tenants TSYS1 routing-instances vr1 interface lt-0/0/0.1
user@SRX# set tenants TSYS1 routing-instances vr1 interface ge-0/0/1.0
```

4. Configure the "security profiles" on page 70. We are only showing the minimal configuration needed to configure logical and tenant systems for this example.

```
[edit]
user@SRX# set system security-profile SP0 logical-system root-ls
user@SRX# set system security-profile SP1 tenant TSYS1
user@SRX# set system security-profile SP2 tenant TSYS2
user@SRX# set system security-profile SP3 tenant TSYS3
```

5. Configure the logical systems. This example using an interconnect VPLS switch requires a logical systems.

   a. Configure the interfaces.

```
[edit]
user@SRX# set logical-systems root-ls interfaces ge-0/0/0 unit 0 family inet address
192.168.10.1/24
user@SRX# set logical-systems root-ls interfaces lt-0/0/0 unit 0 encapsulation ethernet
user@SRX# set logical-systems root-ls interfaces lt-0/0/0 unit 0 peer-unit 100
user@SRX# set logical-systems root-ls interfaces lt-0/0/0 unit 0 family inet address
10.0.1.10/24
```

   b. Configure the static routes.

```
[edit]
user@SRX# set logical-systems root-ls routing-options static route 192.168.1.0/24 next-hop
10.0.1.1
user@SRX# set logical-systems root-ls routing-options static route 192.168.2.0/24 next-hop
10.0.1.2
user@SRX# set logical-systems root-ls routing-options static route 192.168.3.0/24 next-hop
10.0.1.3
```

6. Configure security zones and policies in the logical systems to allow traffic flow from the tenants to the Internet. Additional security policies can be configured on both the logical and tenant systems to allow traffic between tenants.

   a. Configure security zones.

```
[edit]
user@SRX# set logical-systems root-ls security zones security-zone trust host-inbound-
```

```
traffic system-services ping
user@SRX# set logical-systems root-ls security zones security-zone trust interfaces
lt-0/0/0.0
user@SRX# set logical-systems root-ls security zones security-zone untrust host-inbound-
traffic system-services ping
user@SRX# set logical-systems root-ls security zones security-zone untrust interfaces
ge-0/0/0.0
```

b. Configure security policies.

```
[edit]
user@SRX# set logical-systems root-ls security address-book global address TSYS1
192.168.1.0/24
user@SRX# set logical-systems root-ls security address-book global address TSYS2
192.168.2.0/24
user@SRX# set logical-systems root-ls security address-book global address TSYS3
192.168.3.0/24
user@SRX# set logical-systems root-ls security policies from-zone trust to-zone untrust
policy allow-out match source-address TSYS1
user@SRX# set logical-systems root-ls security policies from-zone trust to-zone untrust
policy allow-out match source-address TSYS2
user@SRX# set logical-systems root-ls security policies from-zone trust to-zone untrust
policy allow-out match source-address TSYS3
user@SRX# set logical-systems root-ls security policies from-zone trust to-zone untrust
policy allow-out match destination-address any
user@SRX# set logical-systems root-ls security policies from-zone trust to-zone untrust
policy allow-out match application any
user@SRX# set logical-systems root-ls security policies from-zone trust to-zone untrust
policy allow-out then permit
```

7. Configure security zones and policies in each tenant systems to allow traffic flow to the Internet.

a. Configure security zones.

```
[edit]
user@SRX# set tenants TSYS1 security zones security-zone PC1 host-inbound-traffic system-
services ping
user@SRX# set tenants TSYS1 security zones security-zone PC1 interfaces ge-0/0/1.0
user@SRX# set tenants TSYS1 security zones security-zone VPLS host-inbound-traffic system-
```

```
services ping
user@SRX# set tenants TSYS1 security zones security-zone VPLS interfaces lt-0/0/0.1
```

b. Configure security policies.

```
[edit]
user@SRX# set tenants TSYS1 security address-book global address PC1 192.168.1.0/24
user@SRX# set tenants TSYS1 security policies from-zone PC1 to-zone VPLS policy allow-out
match source-address PC1
user@SRX# set tenants TSYS1 security policies from-zone PC1 to-zone VPLS policy allow-out
match destination-address any
user@SRX# set tenants TSYS1 security policies from-zone PC1 to-zone VPLS policy allow-out
match application any
user@SRX# set tenants TSYS1 security policies from-zone PC1 to-zone VPLS policy allow-out
then permit
```

Results

From configuration mode, confirm your configuration by entering the `show tenants TSYS1` command to verify that the tenant system is created. Enter the `show system login class TSYS1admin1` command to view the permission level for each class that you defined. To ensure that the tenant system administrators are created, enter the `show system login user TSYS1admin1` command. To ensure that the interfaces for interconnect VPLS switch are created, enter the `show interfaces` command. Enter `show logical-systems` to verify the root logical systems configuration.

```
user@SRX# show tenants TSYS1
routing-instances {
    vr1 {
        instance-type virtual-router;
        routing-options {
            static {
                route 0.0.0.0/0 next-hop 10.0.1.10;
            }
        }
        interface lt-0/0/0.1;
        interface ge-0/0/1.0;
    }
}
security {
    address-book {
```

```
        global {
            address PC1 192.168.1.0/24;
        }
    }
    policies {
        from-zone PC1 to-zone VPLS {
            policy allow-out {
                match {
                    source-address PC1;
                    destination-address any;
                    application any;
                }
                then {
                    permit;
                }
            }
        }
    }
    zones {
        security-zone PC1 {
            host-inbound-traffic {
                system-services {
                    ping;
                }
            }
            interfaces {
                ge-0/0/1.0;
            }
        }
        security-zone VPLS {
            host-inbound-traffic {
                system-services {
                    ping;
                }
            }
            interfaces {
                lt-0/0/0.1;
            }
        }
```

```
    }
}
```

```
user@SRX# show system login class TSYS1admin1
tenant TSYS1;
permissions all;
```

```
user@SRX# show system login user TSYS1admin1
uid 2001;
class TSYS1admin1;
authentication {
    encrypted-password "$ABC123";
}
```

```
user@SRX# show interfaces
lt-0/0/0 {
    unit 1 {
        encapsulation ethernet;
        peer-unit 101;
        family inet {
            address 10.0.1.1/24;
        }
    }
    unit 2 {
        encapsulation ethernet;
        peer-unit 102;
        family inet {
            address 10.0.1.2/24;
        }
    }
    unit 3 {
        encapsulation ethernet;
        peer-unit 103;
        family inet {
            address 10.0.1.3/24;
        }
    }
    unit 100 {
        encapsulation ethernet-vpls;
```

```
        peer-unit 0;
    }
    unit 101 {
        encapsulation ethernet-vpls;
        peer-unit 1;
    }
    unit 102 {
        encapsulation ethernet-vpls;
        peer-unit 2;
    }
    unit 103 {
        encapsulation ethernet-vpls;
        peer-unit 3;
    }
}
ge-0/0/1 {
    unit 0 {
        family inet {
            address 192.168.1.254/24;
        }
    }
}
ge-0/0/2 {
    unit 0 {
        family inet {
            address 192.168.2.254/24;
        }
    }
}
ge-0/0/3 {
    unit 0 {
        family inet {
            address 192.168.3.254/24;
        }
    }
}
```

```
user@SRX# show logical-systems
root-ls {
    interfaces {
        ge-0/0/0 {
```

```
            unit 0 {
                family inet {
                    address 192.168.10.1/24;
                }
            }
        }
        lt-0/0/0 {
            unit 0 {
                encapsulation ethernet;
                peer-unit 100;
                family inet {
                    address 10.0.1.10/24;
                }
            }
        }
    }
    routing-options {
        static {
            route 192.168.1.0/24 next-hop 10.0.1.1;
            route 192.168.2.0/24 next-hop 10.0.1.2;
            route 192.168.3.0/24 next-hop 10.0.1.3;
        }
    }
    security {
        address-book {
            global {
                address TSYS1 192.168.1.0/24;
                address TSYS2 192.168.2.0/24;
                address TSYS3 192.168.3.0/24;
            }
        }
        policies {
            from-zone trust to-zone untrust {
                policy allow-out {
                    match {
                        source-address [ TSYS1 TSYS2 TSYS3 ];
                        destination-address any;
                        application any;
                    }
                    then {
                        permit;
                    }
                }
```

```
            }
        }
        zones {
            security-zone trust {
                host-inbound-traffic {
                    system-services {
                        ping;
                    }
                }
                interfaces {
                    lt-0/0/0.0;
                }
            }
            security-zone untrust {
                host-inbound-traffic {
                    system-services {
                        ping;
                    }
                }
                interfaces {
                    ge-0/0/0.0;
                }
            }
        }
    }
}
```

If the output does not display the intended configuration, repeat the configuration instructions in these examples to correct it. If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- Verifying Tenant Systems and Login Configurations Using Primary Administrator | **570**
- Verifying Tenant Systems and Login Configurations Using SSH | **570**
- Verifying PC1 Connectivity to the Internet | **571**

Confirm that the configuration is working properly.

Verifying Tenant Systems and Login Configurations Using Primary Administrator

Purpose

Verify that the tenant systems exist and you can enter them from root as the primary administrator. Return from the tenant system to the root.

Action

From operational mode, use the following command to enter the tenant systems TSYS1:

```
user@SRX> set cli tenant TSYS1
Tenant: TSYS1
user@SRX:TSYS1>
```

Now you are entered to the tenant systems TSYS1. Use the following command to exit from tenant systems TSYS1 to the root:

```
user@SRX:TSYS1> clear cli tenant
Cleared default tenants
user@SRX>
```

Meaning

Tenant system exists and you can enter to the tenant system from the root as the primary administrator.

Verifying Tenant Systems and Login Configurations Using SSH

Purpose

Verify that the tenant systems you created exist, and that the administrator login IDs and passwords that you created are correct.

Action

Use SSH to log in to each user tenant system administrator.

1. Run SSH specifying the IP address of your SRX Series device.

2. Enter the login ID and password for the tenant systems administrator that you created. After you log in, the prompt shows the tenant systems administrator name. Notice how this result differs from the result produced when you log in to the tenant system from the primary logical system at root. Repeat this procedure for all of your tenant systems.

```
login: TSYS1admin1
Password: "$ABC123"

TSYS1admin1@SRX: TSYS1>
```

Meaning

Tenant system administrator `TSYS1admin1` exists and you can login as the tenant system administrator.

Verifying PC1 Connectivity to the Internet

Purpose

Verify end-to-end connectivity.

Action

Ping and run traceroute to the Internet from PC1. In our example the Internet is 192.168.10.254.

1. Run ping from PC1.

```
user@PC1> ping 192.168.10.254 count 2
PING 192.168.10.254 (192.168.10.254): 56 data bytes
64 bytes from 192.168.10.254: icmp_seq=0 ttl=62 time=3.178 ms
64 bytes from 192.168.10.254: icmp_seq=1 ttl=62 time=3.082 ms

--- 192.168.10.254 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 3.082/3.130/3.178/0.048 ms
```

2. Run traceroute from PC1.

```
user@PC1> traceroute 192.168.10.254
traceroute to 192.168.10.254 (192.168.10.254), 30 hops max, 52 byte packets
```

```
1  192.168.1.254 (192.168.1.254)  2.188 ms  1.779 ms  1.896 ms
2  10.0.1.10 (10.0.1.10)  1.888 ms  1.535 ms  1.661 ms
3  192.168.10.254 (192.168.10.254)  3.243 ms  15.077 ms  3.499 ms
```

Meaning

PC1 is able to reach the Internet.

SEE ALSO

Session Creation for Devices Running Tenant Systems | 579

Configuring tenant systems Interconnect with Logical Tunnel Interface point-to-point connection | 596

Release History Table

Release	Description
19.2R1	Starting in Junos OS Release 19.2R1, the virtual-router configured in a tenant system is passed as the default routing-instance to `ping`, `telnet`, `ssh`, `traceroute`, `show arp`, `clear arp`, `show ipv6 neighbors`, and `clear ipv6 neighbors` commands.

RELATED DOCUMENTATION

Logical Systems and Tenant Systems support for VSRX and VSRX 3.0 Instances | 21

Security Zones for Tenant Systems

IN THIS SECTION

- Understanding Zones for Tenant Systems | 573
- Example: Configuring Zones in the Tenant System | 574

Security zones can be configured with tenant systems. For more information see the following topics:

Understanding Zones for Tenant Systems

Security zones are logical entities to which one or more interfaces are bound. Security zones can be configured on the tenant systems by the administrator. On a tenant system, the administrator can configure multiple security zones, dividing the network into network segments to which various security options can be applied.

The primary administrator configures the maximum and reserved numbers of security zones for the tenant system. Then the administrator for the tenant system can create the security zones in the tenant system and assign interfaces to each security zone. The number of zones configured in the tenant system count toward the maximum number of zones available on the device. The `show system security-profile zones` command is used to view the number of security zones allocated to the tenant system and the `show interfaces` command to view the interfaces assigned to the tenant system.

You can configure the following features in a tenant system security zone:

- Interfaces that are part of a security zone.
- Screen options—For every security zone, you can enable a set of predefined screen options that detect and block various kinds of traffic that the device determines as potentially harmful.
- TCP-Reset—When this feature is enabled, the system sends a TCP segment with the RESET flag set when traffic arrives that does not match an existing session and does not have the synchronize flag set.
- Host inbound traffic—This feature specifies the kinds of traffic that can reach the device from systems that are directly connected to its interfaces. You can configure these parameters at the zone level, in which case they affect all interfaces of the zone, or at the interface level. Interface configuration overrides that of the zone.

There are no preconfigured security zones in the tenant system.

The management functional zone (MGT) can be configured for the tenant system. There is the management interface per device that is allocated to the tenant system.

The administrator for the tenant system can configure and view all attributes for a security zone in a tenant system. All security zone attributes in a tenant system are also visible to the primary administrator.

Example: Configuring Zones in the Tenant System

IN THIS SECTION

- Requirements | 574
- Overview | 574
- Configuration | 575
- Verification | 577

This example shows how to configure the zones for the tenant system.

Requirements

Before you begin the configuration:

- Configure the interfaces created by the primary administrator. See *Example: Configuring Interfaces and Routing Instances for a Tenant System*.

Overview

In this example, you can configure zones for the tenant systems. Security zones are the building blocks for policies; they are logical entities to which one or more interfaces are bound. The `[edit tenants tenant-name security zones]` hierarchy level is used to configure the security zones. This example configures the security policies and zones described in Table 34 on page 574.

Table 34: Security Zones Parameters

Feature	Configuration Parameters
Zones 1	• Security zone: trust • System services: any-service • Bind to interfaces xe-0/0/1.0 (trust), xe-0/0/3.0 (untrust)

Table 34: Security Zones Parameters *(Continued)*

Feature	Configuration Parameters
Zone 2	• Security zone: untrust • System services: any-service • Bind to interfaces xe-0/0/1.0 (trust), xe-0/0/3.0 (untrust)

Configuration

IN THIS SECTION

- Procedure | 575

Procedure

CLI Quick Configuration

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

```
set tenants TN1 security zones security-zone trust host-inbound-traffic system-services any-service
set tenants TN1 security zones security-zone trust interfaces xe-0/0/1.0
set tenants TN1 security zones security-zone untrust host-inbound-traffic system-services any-service
set tenants TN1 security zones security-zone untrust interfaces xe-0/0/3.0
```

Step-by-Step Procedure

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

To configure security zones in the tenant system:

1. Define the tenant system name as TN1.

```
[edit]
user@host# set tenants TN1
```

2. Configure a security zone as trust that permits traffic from zone trust and assign it to an interface.

```
[edit tenants TN1 security zones security-zone trust]
user@host# set host-inbound-traffic system-services any-service
user@host# set interfaces xe-0/0/1.0
```

3. Configure a security zone as untrust that permits traffic from zone untrust and assign it to an interface.

```
[edit tenants TN1 security zones security-zone untrust]
user@host# set host-inbound-traffic system-services any-service
user@host# set interfaces xe-0/0/3.0
```

Results

From configuration mode, confirm your configuration by entering the `show tenants tenant-name security policies` and `show tenants tenant-name security zones` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show tenants TN1 security zones
security-zone trust {
    host-inbound-traffic {
        system-services {
            any-service;
        }
    }
    interfaces {
        xe-0/0/1.0;
    }
}
security-zone untrust {
    host-inbound-traffic {
```

```
        system-services {
            any-service;
        }
    }
    interfaces {
        xe-0/0/3.0;
    }
}
```

Verification

IN THIS SECTION

- Verifying Zone Configuration | 577

To confirm that the configuration is working properly, perform the following task:

Verifying Zone Configuration

Purpose

Verify the information about security zones.

Action

To verify the configuration is working properly, enter the `show security zones tenant all` command from operational mode.

```
user@host> show security zones tenant all
```

```
  Tenant: TN1

  Security zone: Host
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
```

```
Interfaces:

Security zone: abc
Send reset for non-SYN session TCP packets: Off
Policy configurable: Yes
Interfaces bound: 0
Interfaces:xe-0/0/1.0

Security zone: def
Send reset for non-SYN session TCP packets: Off
Policy configurable: Yes
Interfaces bound: 1
Interfaces:xe-0/0/3.0
```

Meaning

The output displays the information of security zones configured on the tenant system.

RELATED DOCUMENTATION

Security Policies for Tenant Systems | 654

Flow for Tenant Systems

IN THIS SECTION

- Session Creation for Devices Running Tenant Systems | 579
- Configuring Logical Systems and Tenant Systems Interconnect with Multiple VPLS Switches | 585
- Configuring tenant systems Interconnect with Logical Tunnel Interface point-to-point connection | 596
- Configuring Logical System and Tenant System Interconnect with a Logical Tunnel Interface point-to-point connection | 606

This topic explains how packets are processed in flow sessions on devices that are configured with tenant systems. It describes how the device running tenant systems handles pass-through traffic

between tenant systems. This topic also covers self-traffic as self-initiated traffic within a tenant system and self-traffic terminated on another tenant system. Before addressing tenant systems, the topic provides basic information about the SRX Series architecture with respect to packet processing and sessions. Finally, addresses the sessions and how to change session characteristics.

Session Creation for Devices Running Tenant Systems

IN THIS SECTION

- Understanding Packet Classification | **580**
- Understanding the VPLS Switch and Logical Tunnel Interfaces | **580**
- Handling Pass-Through Traffic for Tenant Systems | **580**
- Handling Self-Traffic | **582**
- Understanding Session and Gate Limitation Control | **584**
- About Configuring Sessions | **584**

A session is created, based on routing and other classification information, to store information and allocate resources for a flow. Basically, a session is established when a traffic enters a tenant system interface, route lookup is performed to identify the next hop interface, and policy lookup is performed.

Optionally, the tenant systems enable you to configure an internal software switch. A virtual private LAN switch (VPLS) is implemented as an interconnect in tenant system. The VPLS enables both transit traffic and traffic terminated at a tenant system to pass between tenant systems. To allow traffic to pass between tenant systems or between tenant system and logical system, logical tunnel (lt-0/0/0) interfaces across the interconnect tenant system are used.

NOTE: Packet sequence occurs at the ingress and the egress interfaces. Packets traversing between tenant systems might not be processed in the order in which they were received on the physical interface.

Understanding Packet Classification

The Packet classification for a flow-based processing is based on both the physical interface and the *logical interface* and depends on the incoming interface. The packet classification is performed at the ingress point and within a flow, the packet-based processing also takes place on an SPU sometimes.

Packet classification is assessed the same way for devices that are configured with or without tenant systems. The traffic for a dedicated interface is classified to the tenant system that contains that interface. The filters and class-of-service features are typically associated with an interface to influence which packets are allowed to transit the device and to apply special actions to packets as needed.

Understanding the VPLS Switch and Logical Tunnel Interfaces

This topic covers the interconnect tenant system that serves as an internal virtual private LAN service (VPLS) switch connecting one tenant system on the device to another. The topic also explains how logical tunnel (lt-0/0/0) interfaces are used to connect tenant systems through the interconnect tenant system.

A device running tenant systems can use an internal VPLS switch to pass traffic without it leaving the device. For communication between tenant systems on the device to occur, you must configure an lt-0/0/0 interface on each tenant system that will use the internal switch, and you must associate it with its peer lt-0/0/0 interface on the interconnect tenant system, effectively creating a logical tunnel between them. You define a peer relationship at each end of the tunnel when you configure the tenant system's lt-0/0/0 interfaces.

You might want all tenant systems on the device to be able to communicate with one another without using an external switch. Alternatively, you might want some tenant systems to connect across the internal switch but not all of them.

WARNING: If you configure an lt-0/0/0 interface in any tenant system and you do not configure a VPLS switch containing a peer lt-0/0/0 interface for it, the commit will fail.

An SRX Series device running tenant systems can be used in a *chassis cluster* and each node has the same configuration.

When you use SRX Series devices configured with tenant systems within a chassis cluster, you must purchase and install the same number of licenses for each node in the chassis cluster. tenant systems licenses pertain to a single chassis, or node, within a chassis cluster and not to the cluster collectively.

Handling Pass-Through Traffic for Tenant Systems

For SRX Series devices running tenant systems, pass-through traffic can exist within a tenant system or between tenant systems.

Pass-Through Traffic Between Tenant Systems

Pass-through traffic between tenant systems is complicated by fact that each tenant system has an ingress and an egress interface that the traffic must transit. It is as if traffic were coming into and going out from two devices. Consider how pass-through traffic is handled between tenant systems given in the topology shown in Figure 14 on page 581.

Figure 14: Tenant Systems, Their Virtual Routers, and Their Interfaces

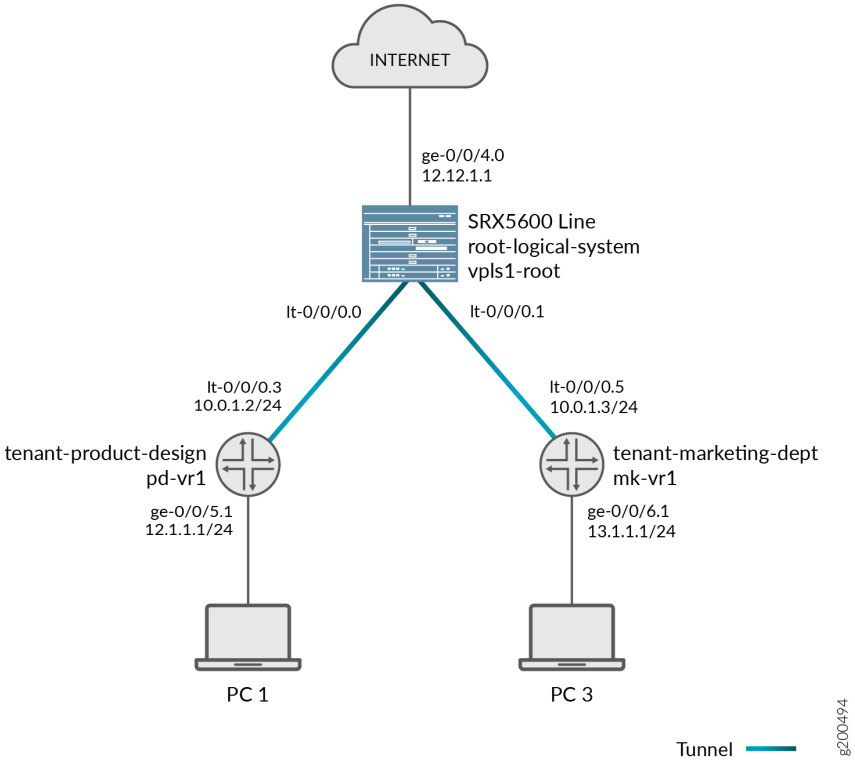

Two sessions must be established for pass-through traffic between tenant systems. (Note that policy lookup is performed in both tenant systems).

- On the incoming tenant system, one session is set up between the ingress interface (a physical interface) and its egress interface (an lt-0/0/0 interface).
- On the egress tenant system, another session is set up between the ingress interface (the lt-0/0/0 interface of the second tenant system) and its egress interface (a physical interface).

Consider how pass-through traffic is handled across tenant systems in the topology shown in Figure 14 on page 581.

- A session is established in the incoming tenant system.
  - When a packet arrives on interface ge-0/0/5, it is identified as belonging to the tenant-product-design tenant system.
  - Because ge-0/0/5 belongs to the pd-vr1 routing instance, route lookup is performed in pd-vr1.
  - As a result of the lookup, the egress interface for the packet is identified as lt-0/0/0.3 with the next hop identified as lt-0/0/0.5, which is the ingress interface in the tenant-marketing-dept.
  - A session is established between ge-0/0/5 and lt-0/0/0.3.
- A session is established in the outgoing tenant system.
  - The packet is injected into the flow again from lt-0/0/0.5, and the tenant system context identified as tenant-marketing-dept is derived from the interface.
  - Packet processing continues in the tenant-marketing-dept tenant system.
  - To identify the egress interface, route lookup for the packet is performed in the mk-vr1 routing instances.
  - The outgoing interface is identified as ge-0/0/6, and the packet is transmitted from the interface to the network.

Handling Self-Traffic

Self-traffic is traffic that originates in a tenant system on a device and is either sent out to the network from that tenant system or is terminated on another tenant system on the device.

Self-Initiated Traffic

Self-initiated traffic is generated from a source tenant system context and forwarded directly to the network from the tenant system interface.

The following process occurs:

- When a packet is generated in a tenant system, a process for handling the traffic is started in the tenant system.
- Route lookup is performed to identify the egress interface, and a session is established.
- The tenant system performs a policy lookup and processes the traffic accordingly.

Consider how self-initiated traffic is handled across tenant systems given the topology shown in Figure 14 on page 581.

- A packet is generated in the tenant-product-design tenant system, and a process for handling the traffic is started in the tenant system.
- Route lookup is performed in pd-vr2, and the egress interface is identified as ge-0/0/8.
- A session is established.
- The packet is transmitted to the network from ge-0/0/8.

Traffic Terminated on a Tenant System

When a packet enters the device on an interface belonging to a tenant system and the packet is destined for another tenant system on the device, the packet is forwarded between the tenant systems in the same manner as is pass-through traffic. However, route lookup in the second tenant system identifies the local egress interface as the packet destination. Consequently the packet is terminated on the second tenant system as self-traffic.

- For terminated self-traffic, two policy lookups are performed, and two sessions are established.
  - On the incoming tenant system, one session is set up between the ingress interface (a physical interface) and its egress interface (an lt-0/0/0 interface).
  - On the destination tenant system, another session is set up between the ingress interface (the lt-0/0/0 interface of the second tenant system) and the local interface.

Consider how terminated self-traffic is handled across tenant systems in the topology shown in Figure 14 on page 581.

- A session is established in the incoming tenant system.
  - When a packet arrives on interface ge-0/0/5, it is identified as belonging to the tenant-product-design tenant system.
  - Because ge-0/0/5 belongs to the pd-vr1 routing instance, route lookup is performed in pd-vr1.
  - As a result of the lookup, the egress interface for the packet is identified as lt-0/0/0.3 with the next hop identified as lt-0/0/0.5, the ingress interface in the ls-marketing-dept.
  - A session is established between ge-0/0/5 and lt-0/0/0.3.
- A management session is established in the destination tenant system.
  - The packet is injected into the flow again from lt-0/0/0.5, and the tenant system context identified as tenant-marketing-dept is derived from the interface.
  - Packet processing continues in the tenant-marketing-dept tenant system.

- Route lookup for the packet is performed in the mk-vr1 routing instance. The packet is terminated in the destination tenant system as self-traffic.

Understanding Session and Gate Limitation Control

Sessions are created based on routing and other classification information to store information and allocate resources for a flow. The tenant systems flow module provides session and gate limitation to ensure that these resources are shared among the tenant systems. Resources allocation and limitations for each tenant system are specified in the security profile bound to the tenant system.

- For session limiting, the system checks the first packet of a session against the maximum number of sessions configured for the tenant system. When the maximum limit of session is reached, the device drops the packet and logs the event.
- For gate limiting, the device checks the first packet of a session against the maximum number of gates configured for the tenant system. If the maximum number of gates for a tenant system is reached, the device rejects the gate open request and logs the event.

About Configuring Sessions

Depending on the protocol and service, a session is programmed with a timeout value. For example, the default timeout for TCP is 1800 seconds. The default timeout for UDP is 60 seconds. When a flow is terminated, it is marked as invalid, and its timeout is reduced to 10 seconds. If no traffic uses the If no traffic uses the session before the service timeout, the session is aged out and freed to a common resource pool for reuse.

You can affect the life of a session in the following ways:

- Age out sessions, based on how full the session table is.
- Set an explicit timeout for aging out TCP sessions.
- Configure a TCP session to be invalidated when it receives a TCP RST (reset) message.
- You can configure sessions to accommodate other systems as follows:
  - Disable TCP packet security checks.
  - Change the maximum segment size.

Configuring Logical Systems and Tenant Systems Interconnect with Multiple VPLS Switches

IN THIS SECTION

- Requirements | 585
- Overview | 585
- Configuration | 586
- Verification | 594

This example shows how to interconnect logical systems and tenant systems with multiple VPLS switches. This is achieved by configuring multiple logical systems and tenant systems with more than one logical tunnel (LT) interface under a tenant system and multiple VPLS switches that are configured to pass the traffic without leaving an SRX Series device.

Requirements

This example uses an SRX Series device running Junos OS with logical systems and tenant systems.

Overview

In this example, we configure multiple LT interfaces and multiple VPLS switches under one tenant system.

In this example, we also configure interconnection between multiple logical systems and tenant systems with LT interface point-to point connections (Encapsulation Ethernet and Encapsulation Frame-Relay).

For interconnected logical systems and tenant systems with multiple VPLS switches, this example configures logical tunnel interfaces lt-0/0/0 with ethernet-vpls as the encapsulation type. The corresponding peer lt-0/0/0 interfaces and security-profiles are assigned to the logical systems and tenant systems. The routing instance for the VPLS switch-1 and VPLS switch-2 are also assigned to the logical systems and tenant systems.

Figure 15 on page 586 shows the topology for interconnected logical systems and tenant systems with multiple VPLS switches.

Figure 15: Configuring the interconnected logical systems and tenant systems with multiple VPLS switches.

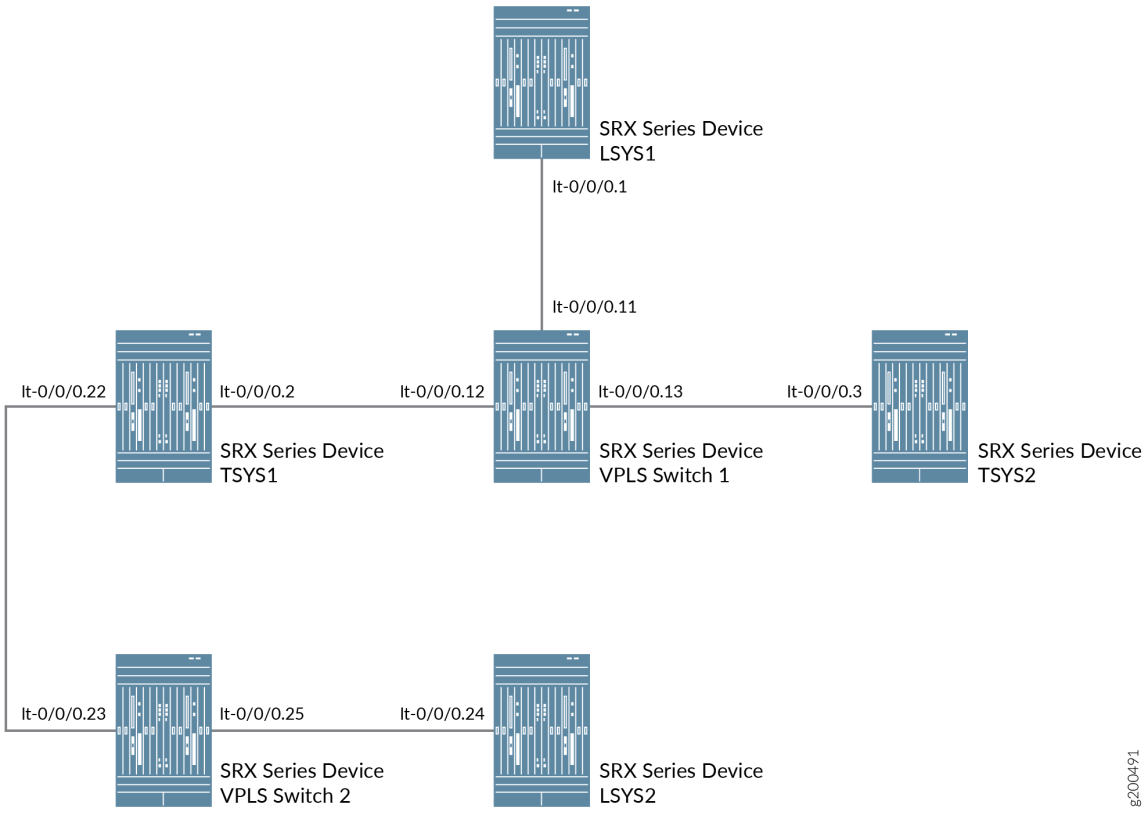

Configuration

IN THIS SECTION

- CLI Quick Configuration | **587**
- Procedure | **588**

To configure interfaces for the logical system and tenant system, perform these tasks:

CLI Quick Configuration

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

```
set interfaces lt-0/0/0 unit 11 encapsulation ethernet-vpls
set interfaces lt-0/0/0 unit 11 peer-unit 1
set interfaces lt-0/0/0 unit 12 encapsulation ethernet-vpls
set interfaces lt-0/0/0 unit 12 peer-unit 2
set interfaces lt-0/0/0 unit 13 encapsulation ethernet-vpls
set interfaces lt-0/0/0 unit 13 peer-unit 3
set interfaces lt-0/0/0 unit 23 encapsulation ethernet-vpls
set interfaces lt-0/0/0 unit 23 peer-unit 22
set interfaces lt-0/0/0 unit 25 encapsulation ethernet-vpls
set interfaces lt-0/0/0 unit 25 peer-unit 24
set routing-instances vpls-switch-1 instance-type vpls
set routing-instances vpls-switch-1 interface lt-0/0/0.11
set routing-instances vpls-switch-1 interface lt-0/0/0.12
set routing-instances vpls-switch-1 interface lt-0/0/0.13
set routing-instances vpls-switch-2 instance-type vpls
set routing-instances vpls-switch-2 interface lt-0/0/0.23
set routing-instances vpls-switch-2 interface lt-0/0/0.25
set logical-systems LSYS1 interfaces lt-0/0/0 unit 1 encapsulation ethernet
set logical-systems LSYS1 interfaces lt-0/0/0 unit 1 peer-unit 11
set logical-systems LSYS1 interfaces lt-0/0/0 unit 1 family inet address 192.168.0.1/24
set interfaces lt-0/0/0 unit 2 encapsulation ethernet
set interfaces lt-0/0/0 unit 2 peer-unit 12
set interfaces lt-0/0/0 unit 2 family inet address 192.168.0.2/24
set interfaces lt-0/0/0 unit 22 encapsulation ethernet
set interfaces lt-0/0/0 unit 22 peer-unit 23
set interfaces lt-0/0/0 unit 22 family inet address 192.168.4.1/30
set tenants TSYS1 routing-instances vr11 instance-type virtual-router
set tenants TSYS1 routing-instances vr11 interface lt-0/0/0.2
set tenants TSYS1 routing-instances vr11 interface lt-0/0/0.22
set interfaces lt-0/0/0 unit 3 encapsulation ethernet
set interfaces lt-0/0/0 unit 3 peer-unit 13
set interfaces lt-0/0/0 unit 3 family inet address 192.168.0.3/24
set tenants TSYS2 routing-instances vr12 instance-type virtual-router
set tenants TSYS2 routing-instances vr12 interface lt-0/0/0.3
set logical-systems LSYS2 interfaces lt-0/0/0 unit 24 encapsulation ethernet
set logical-systems LSYS2 interfaces lt-0/0/0 unit 24 peer-unit 25
```

```
set logical-systems LSYS2 interfaces lt-0/0/0 unit 24 family inet address 192.168.4.2/30
set system security-profile SP-user policy maximum 100
set system security-profile SP-user policy reserved 50
set system security-profile SP-user zone maximum 60
set system security-profile SP-user zone reserved 10
set system security-profile SP-user flow-session maximum 100
set system security-profile SP-user flow-session reserved 50
set system security-profile SP-user logical-system LSYS1
set system security-profile SP-user tenant TSYS1
set system security-profile SP-user tenant TSYS2
set system security-profile SP-user logical-system LSYS2
```

Procedure

Step-by-Step Procedure

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

1. Configure the lt-0/0/0 interfaces.

```
[edit]
user@host# set interfaces lt-0/0/0 unit 11 encapsulation ethernet-vpls
user@host# set interfaces lt-0/0/0 unit 11 peer-unit 1
user@host# set interfaces lt-0/0/0 unit 12 encapsulation ethernet-vpls
user@host# set interfaces lt-0/0/0 unit 12 peer-unit 2
user@host# set interfaces lt-0/0/0 unit 13 encapsulation ethernet-vpls
user@host# set interfaces lt-0/0/0 unit 13 peer-unit 3
user@host# set interfaces lt-0/0/0 unit 23 encapsulation ethernet-vpls
user@host# set interfaces lt-0/0/0 unit 23 peer-unit 22
user@host# set interfaces lt-0/0/0 unit 25 encapsulation ethernet-vpls
user@host# set interfaces lt-0/0/0 unit 25 peer-unit 24
```

2. Configure the routing instance for the VPLS switches and add interfaces to it.

```
[edit]
user@host# set routing-instances vpls-switch-1 instance-type vpls
user@host# set routing-instances vpls-switch-1 interface lt-0/0/0.11
user@host# set routing-instances vpls-switch-1 interface lt-0/0/0.12
user@host# set routing-instances vpls-switch-1 interface lt-0/0/0.13
```

```
user@host# set routing-instances vpls-switch-2 instance-type vpls
user@host# set routing-instances vpls-switch-2 interface lt-0/0/0.23
user@host# set routing-instances vpls-switch-2 interface lt-0/0/0.25
```

3. Configure LSYS1 with lt-0/0/0.1 interface and peer lt-0/0/0.11.

```
[edit]
user@host# set logical-systems LSYS1 interfaces lt-0/0/0 unit 1 encapsulation ethernet
user@host# set logical-systems LSYS1 interfaces lt-0/0/0 unit 1 peer-unit 11
user@host# set logical-systems LSYS1 interfaces lt-0/0/0 unit 1 family inet address
192.168.0.1/24
```

4. Configure TSYS1 with lt-0/0/0.2 interface and peer lt-0/0/0.12.

```
[edit]
user@host# set interfaces lt-0/0/0 unit 2 encapsulation ethernet
user@host# set interfaces lt-0/0/0 unit 2 peer-unit 12
user@host# set interfaces lt-0/0/0 unit 2 family inet address 192.168.0.2/24
user@host# set interfaces lt-0/0/0 unit 22 encapsulation ethernet
user@host# set interfaces lt-0/0/0 unit 22 peer-unit 23
user@host# set interfaces lt-0/0/0 unit 22 family inet address 192.168.4.1/30
user@host# set tenants TSYS1 routing-instances vr11 instance-type virtual-router
user@host# set tenants TSYS1 routing-instances vr11 interface lt-0/0/0.2
user@host# set tenants TSYS1 routing-instances vr11 interface lt-0/0/0.22
```

5. Configure TSYS2 with lt-0/0/0.3 interface and peer lt-0/0/0.13

```
[edit]
user@host# set interfaces lt-0/0/0 unit 3 encapsulation ethernet
user@host# set interfaces lt-0/0/0 unit 3 peer-unit 13
user@host# set interfaces lt-0/0/0 unit 3 family inet address 192.168.0.3/24
user@host# set tenants TSYS2 routing-instances vr12 instance-type virtual-router
user@host# set tenants TSYS2 routing-instances vr12 interface lt-0/0/0.3
```

6. Configure LSYS2 with lt-0/0/0 interface and peer-unit 24.

```
[edit]
user@host# set logical-systems LSYS2 interfaces lt-0/0/0 unit 24 encapsulation ethernet
user@host# set logical-systems LSYS2 interfaces lt-0/0/0 unit 24 peer-unit 25
```

```
user@host# set logical-systems LSYS2 interfaces lt-0/0/0 unit 24 family inet address
192.168.4.2/30
```

7. Assign security-profile for logical-systems.

```
[edit]
user@host# set system security-profile SP-user policy maximum 100
user@host# set system security-profile SP-user policy reserved 50
user@host# set system security-profile SP-user zone maximum 60
user@host# set system security-profile SP-user zone reserved 10
user@host# set system security-profile SP-user flow-session maximum 100
user@host#set system security-profile SP-user flow-session reserved 50
user@host# set system security-profile SP-user logical-system LSYS1
user@host# set system security-profile SP-user tenant TSYS1
user@host# set system security-profile SP-user tenant TSYS2
user@host# set system security-profile SP-user logical-system LSYS2
```

Results

- From configuration mode, confirm your configuration by entering the `show interfaces lt-0/0/0`, command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it

```
unit 2 {
    encapsulation ethernet;
    peer-unit 12;
    family inet {
        address 192.168.0.2/24;
    }
}
unit 3 {
    encapsulation ethernet;
    peer-unit 13;
    family inet {
        address 192.168.0.3/24;
    }
}
unit 11 {
    encapsulation ethernet-vpls;
    peer-unit 1;
```

```
}
unit 12 {
    encapsulation ethernet-vpls;
    peer-unit 2;
}
unit 13 {
    encapsulation ethernet-vpls;
    peer-unit 3;
}
unit 22 {
    encapsulation ethernet;
    peer-unit 23;
    family inet {
        address 192.168.4.1/30;
    }
}
unit 23 {
    encapsulation ethernet-vpls;
    peer-unit 22;
}
unit 25 {
    encapsulation ethernet-vpls;
    peer-unit 24;
}
```

- From configuration mode, confirm your configuration by entering the `show routing-instances`, command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show routing-instances
vpls-switch-1 {
    instance-type vpls;
    interface lt-0/0/0.11;
    interface lt-0/0/0.12;
    interface lt-0/0/0.13;
}
vpls-switch-2 {
    instance-type vpls;
    interface lt-0/0/0.23;
```

```
    interface lt-0/0/0.25;
}
```

- From configuration mode, confirm your configuration by entering the `show logical-systems LSYS1`, command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show logical-systems LSYS1
interfaces {
    lt-0/0/0 {
        unit 1 {
            encapsulation ethernet;
            peer-unit 11;
            family inet {
                address 192.168.0.1/24;
            }
        }
    }
}
```

- From configuration mode, confirm your configuration by entering the `show logical-systems LSYS2`, command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show tenants TSYS1
routing-instances {
    vr11 {
        instance-type virtual-router;
        interface lt-0/0/0.2;
        interface lt-0/0/0.22;
    }
}
```

- From configuration mode, confirm your configuration by entering the `show logical-systems LSYS3`, command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show tenants TSYS2
routing-instances {
    vr12 {
        instance-type virtual-router;
        interface lt-0/0/0.3;
    }
}
```

- From configuration mode, confirm your configuration by entering the `show logical-systems LSYS2`, command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show logical-systems LSYS2
interfaces {
    lt-0/0/0 {
        unit 24 {
            encapsulation ethernet;
            peer-unit 25;
            family inet {
                address 192.168.4.2/30;
            }
        }
    }
}
```

- From configuration mode, confirm your configuration by entering the `show system security-profile`, command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show system security-profile
SP-user {
    policy {
        maximum 100;
```

```
        reserved 50;
    }
    zone {
        maximum 60;
        reserved 10;
    }
    flow-session {
        maximum 100;
        reserved 50;
    }
    logical-system [ LSYS1 LSYS2 ];
    tenant [ TSYS1 TSYS2 ];
}
```

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

Verification

IN THIS SECTION

- Verifying the Security-Profile for Logical-systems | **594**
- Verifying the LT Interfaces for Logical systems | **595**

To confirm that the configuration is working properly, perform these tasks:

Verifying the Security-Profile for Logical-systems

Purpose

Verify security profile for each logical systems.

Action

From operational mode, enter the `show system security-profile security-log-stream-number logical-system all` command.

```
user@host> show system security-profile assignment summary
```

```
                                  Total     Maximum
security-profiles                     1          65
logical-systems                       1          32
tenants                               0          32
logical-systems and tenants           1          64
```

Meaning

The output provides the usage and reserved values for the logical systems when security-log-stream is configured.

Verifying the LT Interfaces for Logical systems

Purpose

Verify interfaces for logical systems.

Action

From operational mode, enter the `show interfaces lt-0/0/0 terse` command.

```
user@host> show interfaces lt-0/0/0 terse
```

```
Interface               Admin Link Proto    Local                 Remote
lt-0/0/0                up    up
lt-0/0/0.1              up    up   inet     192.168.0.1/24
lt-0/0/0.2              up    up   inet     192.168.0.2/24
lt-0/0/0.3              up    up   inet     192.168.0.3/24
lt-0/0/0.11             up    up   vpls
lt-0/0/0.12             up    up   vpls
```

```
lt-0/0/0.13             up    up   vpls
lt-0/0/0.22             up    up   inet     192.168.4.1/30
lt-0/0/0.23             up    up   vpls
lt-0/0/0.24             up    up   inet     192.168.4.2/30
lt-0/0/0.25             up    up   vpls
lt-0/0/0.32767          up    up
```

Meaning

The output provides the status of LT interfaces. All the LT interfaces are up.

Configuring tenant systems Interconnect with Logical Tunnel Interface point-to-point connection

IN THIS SECTION

- Requirements | 596
- Overview | 596
- Configuration | 597
- Verification | 605

This example shows how to interconnect tenant systems with logical tunnel (LT) interfaces in a point-to-point connection.

Requirements

This example uses an SRX Series device running Junos OS with logical systems and tenant systems.

Overview

In this example we show how to interconnect tenant systems with logical tunnel (LT) interface in a point-to-point connection.

For the interconnected tenant systems with a point-to-point connection (encapsulation frame-relay) LT interface, this example configures the logical tunnel interface lt-0/0/0. This example configures security-zone and assigns interfaces to the logical systems.

The interconnected logical system lt-0/0/0 interface is configured with frame-relay as the encapsulation type. The corresponding peer lt-0/0/0 interfaces in the tenant systems are configured with frame-relay as the encapsulation type. A security profile is assigned to the tenant systems.

Figure 16 on page 597 shows the topology for interconnected tenant systems with a point-to-point connection LT interface.

Figure 16: Configuring the interconnect tenant systems with a point-to-point connection LT interface

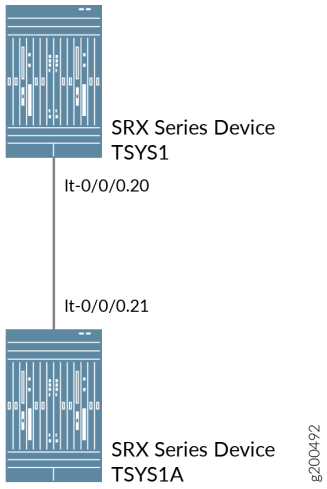

Configuration

IN THIS SECTION

- CLI Quick Configuration | **598**
- Configuring [item] | **599**
- Results | **602**

To configure security-zone and assigns interfaces to tenant systems, perform these tasks:

CLI Quick Configuration

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

```
set system security-profile sp1 tenant TSYS1
set system security-profile sp2 tenant TSYS1A
set interfaces xe-0/0/5 gigether-options redundant-parent reth0
set interfaces xe-0/0/6 gigether-options redundant-parent reth1
set interfaces xe-1/0/5 gigether-options redundant-parent reth0
set interfaces xe-1/0/6 gigether-options redundant-parent reth1
set interfaces reth0 redundant-ether-options redundancy-group 2
set interfaces reth1 redundant-ether-options redundancy-group 1
set interfaces lt-0/0/0 unit 20 encapsulation ethernet
set interfaces lt-0/0/0 unit 20 peer-unit 21
set interfaces lt-0/0/0 unit 20 family inet address 198.51.1.20/24
set interfaces reth0 unit 0 family inet address 198.51.100.1/24
set interfaces lt-0/0/0 unit 21 encapsulation ethernet
set interfaces lt-0/0/0 unit 21 peer-unit 20
set interfaces lt-0/0/0 unit 21 family inet address 198.51.1.21/24
set interfaces reth1 unit 0 family inet address 192.0.2.1/24
set tenants TSYS1 routing-instances vr11 instance-type virtual-router
set tenants TSYS1 routing-instances vr11 interface lt-0/0/0.20
set tenants TSYS1 routing-instances vr11 interface reth0.0
set tenants TSYS1 routing-instances vr11 routing-options static route 192.0.2.0/24 next-hop
198.51.1.21
set tenants TSYS1 security policies default-policy permit-all
set tenants TSYS1 security zones security-zone trust host-inbound-traffic system-services all
set tenants TSYS1 security zones security-zone trust host-inbound-traffic protocols all
set tenants TSYS1 security zones security-zone trust interfaces reth0.0
set tenants TSYS1 security zones security-zone untrust host-inbound-traffic system-services all
set tenants TSYS1 security zones security-zone untrust host-inbound-traffic protocols all
set tenants TSYS1 security zones security-zone untrust interfaces lt-0/0/0.20
set tenants TSYS1A routing-instances vr12 instance-type virtual-router
set tenants TSYS1A routing-instances vr12 interface lt-0/0/0.21
set tenants TSYS1A routing-instances vr12 interface reth1.0
set tenants TSYS1A routing-instances vr12 routing-options static route 198.51.100.0/24 next-hop
198.51.1.20
set tenants TSYS1A security policies default-policy permit-all
set tenants TSYS1A security zones security-zone trust host-inbound-traffic system-services all
set tenants TSYS1A security zones security-zone trust host-inbound-traffic protocols all
```

```
set tenants TSYS1A security zones security-zone trust interfaces reth1.0
set tenants TSYS1A security zones security-zone untrust host-inbound-traffic system-services all
set tenants TSYS1A security zones security-zone untrust host-inbound-traffic protocols all
set tenants TSYS1A security zones security-zone untrust interfaces lt-0/0/0.21
```

Configuring [item]

Step-by-Step Procedure

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

1. Define a security profile sp1 and assign to a tenant system TNI. Define another security profile sp1 and assign to a tenant system TSYS1A

```
[edit]
user@host# set system security-profile sp1 tenant TSYS1
user@host# set system security-profile sp2 tenant TSYS1A
```

2. Set the interface for reth0 and reth1 and assign it to the redundancy group 1 and redundancy group 2.

```
[edit]
set interfaces xe-0/0/5 gigether-options redundant-parent reth0
set interfaces xe-0/0/6 gigether-options redundant-parent reth1
set interfaces xe-1/0/5 gigether-options redundant-parent reth0
set interfaces xe-1/0/6 gigether-options redundant-parent reth1
set interfaces reth0 redundant-ether-options redundancy-group 2
set interfaces reth1 redundant-ether-options redundancy-group 1
```

3. Set the LT interface as encapsulation ethernet in the tenant system TSYS1.

```
[edit]
user@host# set interfaces lt-0/0/0 unit 20 encapsulation ethernet
```

4. Configure a peer unit relationship between LT interfaces, thus creating a point-to-point connection.

```
[edit]
user@host# set interfaces lt-0/0/0 unit 20 peer-unit 21
```

5. Specify the IP address for the LT interface.

```
[edit]
user@host# set interfaces lt-0/0/0 unit 20 family inet address 198.51.1.20/24
```

6. Specify the IP address for the reth0.

```
[edit]
user@host# set interfaces reth0 unit 0 family inet address 198.51.100.1/24
```

7. Set the LT interface as encapsulation ethernet in the tenant system TSYS1A.

```
[edit]
user@host# set interfaces lt-0/0/0 unit 21 encapsulation ethernet
```

8. Configure a peer unit relationship between LT interfaces, thus creating a point-to-point connection.

```
[edit]
user@host# set interfaces lt-0/0/0 unit 21 peer-unit 20
```

9. Specify the IP address for the LT interface.

```
[edit]
user@host# set interfaces lt-0/0/0 unit 21 family inet address 198.51.1.21/24
```

10. Specify the IP address for the reth1.

```
[edit]
user@host# set interfaces reth1 unit 0 family inet address 192.0.2.1/24
```

11. Define the routing-instances for TSYS1.

```
[edit]
set tenants TSYS1 routing-instances vr11 instance-type virtual-router
set tenants TSYS1 routing-instances vr11 interface lt-0/0/0.20
set tenants TSYS1 routing-instances vr11 interface reth0.0
set tenants TSYS1 routing-instances vr11 routing-options static route 192.0.2.0/24 next-hop
198.51.1.21
```

12. Configure a security policy that permits all traffics.

```
[edit]
user@host# set tenants TSYS1 security policies default-policy permit-all
```

13. Configure security zones.

```
[edit]
set tenants TSYS1 security zones security-zone trust host-inbound-traffic system-services
all
set tenants TSYS1 security zones security-zone trust host-inbound-traffic protocols all
set tenants TSYS1 security zones security-zone trust interfaces reth0.0
set tenants TSYS1 security zones security-zone untrust host-inbound-traffic system-services
all
set tenants TSYS1 security zones security-zone untrust host-inbound-traffic protocols all
set tenants TSYS1 security zones security-zone untrust interfaces lt-0/0/0.20
```

14. Define the routing-instances for TSYS1A.

```
[edit]
set tenants TSYS1A routing-instances vr12 instance-type virtual-router
set tenants TSYS1A routing-instances vr12 interface lt-0/0/0.21
set tenants TSYS1A routing-instances vr12 interface reth1.0
set tenants TSYS1A routing-instances vr12 routing-options static route 198.51.100.0/24 next-
hop 198.51.1.20
```

15. Configure a security policy that permits all traffics.

```
[edit]
set tenants TSYS1A security policies default-policy permit-all
```

16. Configure security zones.

```
[edit]
set tenants TSYS1A security zones security-zone trust host-inbound-traffic system-services
all
set tenants TSYS1A security zones security-zone trust host-inbound-traffic protocols all
set tenants TSYS1A security zones security-zone trust interfaces reth1.0
set tenants TSYS1A security zones security-zone untrust host-inbound-traffic system-
services all
set tenants TSYS1A security zones security-zone untrust host-inbound-traffic protocols all
set tenants TSYS1A security zones security-zone untrust interfaces lt-0/0/0.21
```

Results

- From configuration mode, confirm your configuration by entering the `show tenants TSYS1` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show tenants TSYS1
routing-instances {
    vr11 {
        instance-type virtual-router;
        interface lt-0/0/0.20;
        interface reth0.0;
        routing-options {
            static {
                route 192.0.2.0/24 next-hop 198.51.1.21;
            }
        }
    }
}
security {
    policies {
        default-policy {
```

```
            permit-all;
        }
    }
    zones {
        security-zone trust {
            host-inbound-traffic {
                system-services {
                    all;
                }
                protocols {
                    all;
                }
            }
            interfaces {
                reth0.0;
            }
        }
        security-zone untrust {
            host-inbound-traffic {
                system-services {
                    all;
                }
                protocols {
                    all;
                }
            }
            interfaces {
                lt-0/0/0.20;
            }
        }
    }
}
```

- From configuration mode, confirm your configuration by entering the `show tenants TSYS1A` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show tenants TSYS1A
routing-instances {
    vr12 {
        instance-type virtual-router;
```

```
        interface lt-0/0/0.21;
        interface reth1.0;
        routing-options {
            static {
                route 198.51.100.0/24 next-hop 198.51.1.20;
            }
        }
    }
}
security {
    policies {
        default-policy {
            permit-all;
        }
    }
    zones {
        security-zone trust {
            host-inbound-traffic {
                system-services {
                    all;
                }
                protocols {
                    all;
                }
            }
            interfaces {
                reth1.0;
            }
        }
        security-zone untrust {
            host-inbound-traffic {
                system-services {
                    all;
                }
                protocols {
                    all;
                }
            }
            interfaces {
                lt-0/0/0.21;
            }
        }
```

```
    }
}
```

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

Verification

IN THIS SECTION

- Verifying the Security-Profile for all tenant systems | **605**

To confirm that the configuration is working properly, perform these tasks:

Verifying the Security-Profile for all tenant systems

Purpose

Verify security profile for each logical systems.

Action

From operational mode, enter the `show system security-profile zone tenant al` command.

```
user@host> show system security-profile zone tenant al
```

```
logical-system tenant name   security profile name       usage    reserved     maximum

T1                           bronze                        1           0        2048
T1A                            pX                            0           0        2048
```

Meaning

The output provides the usage and reserved values for the logical systems when security-log-stream is configured.

Configuring Logical System and Tenant System Interconnect with a Logical Tunnel Interface point-to-point connection

IN THIS SECTION

- Requirements | 606
- Overview | 606
- Configuration | 607
- Verification | 612

This example shows how to interconnect logical systems and tenant systems with logical tunnel (LT) interface in a point-to-point connection.

Requirements

This example uses an SRX Series device running Junos OS with logical systems and tenant systems.

Overview

In this example we show how to interconnect logical systems and tenant systems with logical tunnel (LT) interface point-to-point connection.

For the interconnect logical system and tenant system with a point-to-point connection LT interface, the example configures logical tunnel interfaces lt-0/0/0. This example configures security-zone and assigns interfaces to the logical systems

To interconnect the logical system and tenant system, lt-0/0/0 interfaces are configured with Ethernet as the encapsulation type. The corresponding peer lt-0/0/0 interfaces are configured with Ethernet as the encapsulation type. A security profile is assigned to the logical system and tenant system

Figure 17 on page 607 shows the topology for interconnected logical systems and tenant systems with LT interface point-to-point connection.

Figure 17: Configuring the interconnect between logical systems and tenant systems with a point-to-point connection LT interface

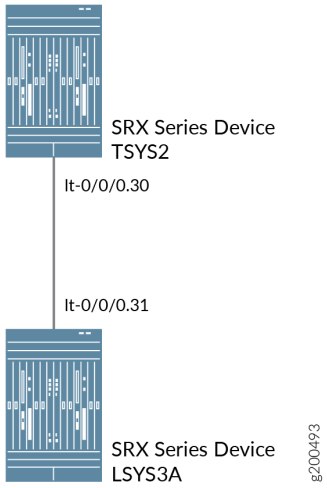

Configuration

IN THIS SECTION

- CLI Quick Configuration | **607**
- Procedure | **608**
- Results | **611**

To configure security-zone and assigns interfaces to logical systems, perform these tasks:

CLI Quick Configuration

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

```
set system security-profile SP-user tenant TSYS2
set interfaces lt-0/0/0 unit 30 encapsulation ethernet
set interfaces lt-0/0/0 unit 30 peer-unit 31
set interfaces lt-0/0/0 unit 30 family inet address 192.255.2.1/30
set tenants TSYS2 routing-instances vr11 instance-type virtual-router
```

```
set tenants TSYS2 routing-instances vr11 interface lt-0/0/0.30
set security zones security-zone LT interfaces lt-0/0/0.30
set system security-profile SP-user logical-system LSYS3A
set logical-systems LSYS3A interfaces lt-0/0/0 unit 21 encapsulation ethernet
set logical-systems LSYS3A interfaces lt-0/0/0 unit 21 peer-unit 20
set logical-systems LSYS3A interfaces lt-0/0/0 unit 21 family inet address 192.255.2.2/30
set logical-systems LSYS3A security policies from-zone LT to-zone LT policy LT match source-
address any
set logical-systems LSYS3A security policies from-zone LT to-zone LT policy LT match destination-
address any
set logical-systems LSYS3A security policies from-zone LT to-zone LT policy LT match application
any
set logical-systems LSYS3A security policies from-zone LT to-zone LT policy LT then permit
set logical-systems LSYS3A security policies default-policy permit-all
set logical-systems LSYS3A security zones security-zone LT host-inbound-traffic system-services
all
set logical-systems LSYS3A security zones security-zone LT host-inbound-traffic protocols all
set logical-systems LSYS3A security zones security-zone LT interfaces lt-0/0/0.31
```

Procedure

Step-by-Step Procedure

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

1. Define a security profile and assign to a tenant system.

```
[edit]
user@host# set system security-profile SP-user tenant TSYS2
```

2. Set the LT interface as encapsulation ethernet in the tenant system.

```
[edit]
user@host# set interfaces lt-0/0/0 unit 20 encapsulation ethernet
```

3. Configure a peer relationship for tenant systems TSYS2.

```
[edit]
user@host# set interfaces lt-0/0/0 unit 20 peer-unit 21
```

4. Specify the IP address for the LT interface.

```
[edit]
user@host# set interfaces lt-0/0/0 unit 20 family inet address 192.255.2.1/30
```

5. Set the security zone for the LT interface.

```
[edit]
user@host# set logical-systems LSYS2 security zones security-zone LT interfaces lt-0/0/0.30
```

6. Define a security profile and assign to a logical system.

```
[edit]
user@host# set system security-profile SP-user logical-system LSYS3A
```

7. Define the routing-instances for TSYS2.

```
[edit]
set tenants TSYS2 routing-instances vr11 instance-type virtual-router
set tenants TSYS2 routing-instances vr11 interface lt-0/0/0.30
```

8. Set the LT interface as encapsulation ethernet in the logical system 3A.

```
[edit]
user@host# set logical-systems LSYS3A interfaces lt-0/0/0 unit 21 encapsulation ethernet
```

9. Configure a peer relationship for logical systems LSYS3A.

```
[edit]
user@host# set logical-systems LSYS3A interfaces lt-0/0/0 unit 21 peer-unit 20
```

10. Specify the IP address for the LT interface.

```
[edit]
user@host# set logical-systems LSYS3A interfaces lt-0/0/0 unit 21 family inet address
192.255.2.2/30
```

11. Configure a security policy that permits traffic from the LT zone to the LT policy LT zone.

```
[edit]
user@host# set logical-systems LSYS3A security policies from-zone LT to-zone LT policy LT
match source-address any
user@host# set logical-systems LSYS3A security policies from-zone LT to-zone LT policy LT
match destination-address any
user@host# set logical-systems LSYS3A security policies from-zone LT to-zone LT policy LT
match application any
user@host# set logical-systems LSYS3A security policies from-zone LT to-zone LT policy LT
then permit
```

12. Configure a security policy that permits traffic from default-policy.

```
[edit]
user@host# set logical-systems LSYS3A security policies default-policy permit-all
```

13. Configure security zones.

```
[edit]
user@host# set logical-systems LSYS3A security zones security-zone LT host-inbound-traffic
system-services all
user@host# set logical-systems LSYS3A security zones security-zone LT host-inbound-traffic
protocols all
user@host# set logical-systems LSYS3A security zones security-zone LT interfaces
lt-0/0/0.31
```

Results

- From configuration mode, confirm your configuration by entering the `show tenants TSYS2` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show tenants TSYS2
routing-instances {
    vr11 {
        instance-type virtual-router;
        interface lt-0/0/0.30;
    }
}
```

- From configuration mode, confirm your configuration by entering the `show logical-systems LSYS3A` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show logical-systems LSYS3A
interfaces {
    lt-0/0/0 {
        unit 21 {
            encapsulation ethernet;
            peer-unit 20;
            family inet {
                address 192.255.2.2/30;
            }
        }
    }
}
security {
    policies {
        from-zone LT to-zone LT {
            policy LT {
                match {
                    source-address any;
                    destination-address any;
                    application any;
                }
```

```
                then {
                    permit;
                }
            }
        }
        default-policy {
            permit-all;
        }
    }
    zones {
        security-zone LT {
            host-inbound-traffic {
                system-services {
                    all;
                }
                protocols {
                    all;
                }
            }
            interfaces {
                lt-0/0/0.31;
            }
        }
    }
}
```

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

Verification

IN THIS SECTION

- Verifying the LT Interfaces for all Logical and tenant systems | **613**
- Verifying the Security-Profile for all Logical-systems | **613**

To confirm that the configuration is working properly, perform these tasks:

Verifying the LT Interfaces for all Logical and tenant systems

Purpose

Verify interfaces for logical systems.

Action

From operational mode, enter the `show system security-profile zone all-logical-systems-tenants` command.

```
user@host> show system security-profile zone all-logical-systems-tenants
```

```
logical-system tenant name   security profile name       usage    reserved     maximum

root-logical-system          Default-Profile                1           0        2048
LSYS3A1                        gold                           1           0        2048
TSYS23                          bronze                         1           0        2048
```

Meaning

The output provides the status of LT interfaces. All the LT interfaces are up.

Verifying the Security-Profile for all Logical-systems

Purpose

Verify security profile for each logical systems.

Action

From operational mode, enter the `show system security-profile security-log-stream-number logical-system all` command.

```
user@host> show system security-profile security-log-stream-number logical-system all
```

```
logical system name   security profile name       usage    reserved   maximum
root-logical-system   Default-Profile                2           0          2000
LSYS3A                 SP-user                        1          10          60
```

Meaning

The output provides the usage and reserved values for the logical systems when security-log-stream is configured.

Flow Trace for Tenant Systems

IN THIS SECTION

- Flow Trace Support for Tenant Systems Overview | **615**
- Configure Flow Trace Support for Tenant Systems | **615**

Flow trace also called traceoptions, allows you to monitor traffic flow into and out of an SRX Series device. You can use tracoptions as debugging tool to trace the packets as they traverse the SRX Series device. Traceoptions help you to get details of actions by your security device.

Flow Trace Support for Tenant Systems Overview

For an SRX Series device configured with tenant systems, by default the traceoptions are configured at the root level only. In this case, all the system traces including root and tenant systems are logged in one single trace file. This generated large amounts of information in a single file.

Starting in Junos OS Release 19.4R1, you can enable tracing operations per tenant system level. When you configure the traceoptions at the tenant system level, then the traces for that specific tenant systems are logged in the respective trace file. You can generate an output file for the specified tenant system, and you can find the required traffic information easily in the trace file.

When you enable traceoptions, you specify the name of the file and the type of information you want to trace.

All flow trace sent to one log file in root, if you enable the traceoptions under root context. Traces for a tenant system only sent to the respective trace file, if you enable the traceoptions for the specific tenant system.

Configure Flow Trace Support for Tenant Systems

Configuring traceoptions for a tenant system includes configuring both a target file and a flag. The target file determines where the trace output is recorded. The flag defines what type of data to be collected. If you configure traceoptions for a tenant system, the respective trace file sent to the specific tenant system log file only.

To configure traceoptions for a tenant system:

1. Create tenant system `TSYS1` and setup the basic configurations. See "Tenant System Configuration Overview" on page 537.
2. Configure target file to save the trace information for the tenant system.

```
[edit]
user@host# set tenants TSYS1 security flow traceoptions file flow_tsys1.log
user@host# set tenants TSYS1 security flow traceoptions file size 1g
```

3. Configure traceoptions flag for the tenant system.

```
[edit]
user@host# set tenants TSYS1 security flow traceoptions flag all
```

After you commit the traceoptions configuration, you can view the traceoptions debug files for the tenant system using `show log` *`tracefilename`* operational command.

```
user@host:TSYS1> show log flow_tsys1.logNov  7 13:21:47
13:21:47.217744:CID-0:THREAD_ID-05:LSYS_ID-32:RT:<192.0.2.0/0->198.51.100.0/9011;1,0x0> :

Nov  7 13:21:47 13:21:47.217747:CID-0:THREAD_ID-05:LSYS_ID-32:RT:packet [84] ipid = 39281,
@0x7f490ae56d52

Nov  7 13:21:47 13:21:47.217749:CID-0:THREAD_ID-05:LSYS_ID-32:RT:---- flow_process_pkt: (thd 5):
flow_ctxt type 0, common flag 0x0, mbuf 0x4882b600, rtbl7

Nov  7 13:21:47 13:21:47.217752:CID-0:THREAD_ID-05:LSYS_ID-32:RT: flow process pak fast ifl 88
in_ifp lt-0/0/0.101

Nov  7 13:21:47 13:21:47.217753:CID-0:THREAD_ID-05:LSYS_ID-32:RT:  lt-0/0/0.101:192.0.2.0-
>198.51.100.0, icmp, (0/0)

Nov  7 13:21:47 13:21:47.217756:CID-0:THREAD_ID-05:LSYS_ID-32:RT: find flow: table 0x11d0a2680,
hash 20069(0xffff), sa 192.0.2.0, da 198.51.100.0, sp 0, d0

Nov  7 13:21:47 13:21:47.217760:CID-0:THREAD_ID-05:LSYS_ID-32:RT:Found: session id 0x12. sess
tok 28685

Nov  7 13:21:47 13:21:47.217761:CID-0:THREAD_ID-05:LSYS_ID-32:RT:  flow got session.

Nov  7 13:21:47 13:21:47.217761:CID-0:THREAD_ID-05:LSYS_ID-32:RT:  flow session id 18

Nov  7 13:21:47 13:21:47.217763:CID-0:THREAD_ID-05:LSYS_ID-32:RT: vector bits 0x200 vector
0x84ae85f0

Nov  7 13:21:47 13:21:47.217764:CID-0:THREAD_ID-05:LSYS_ID-32:RT:set nat 0x11e463550(18) timeout
const to 2

Nov  7 13:21:47 13:21:47.217765:CID-0:THREAD_ID-05:LSYS_ID-32:RT: set_nat_timeout 2 on session 18

Nov  7 13:21:47 13:21:47.217765:CID-0:THREAD_ID-05:LSYS_ID-32:RT:refresh nat 0x11e463550(18)
timeout to 2

Nov  7 13:21:47 13:21:47.217767:CID-0:THREAD_ID-05:LSYS_ID-32:RT:insert usp tag for apps
```

```
Nov  7 13:21:47 13:21:47.217768:CID-0:THREAD_ID-05:LSYS_ID-32:RT:mbuf 0x4882b600, exit nh
0xfffb0006
```

Firewall Authentication for Tenant Systems

IN THIS SECTION

- Understanding Tenant System Firewall Authentication | **617**
- Configuring Firewall Authentication for a Tenant System | **620**
- Understanding Integrated User Firewall Support in a Tenant System | **635**
- Example: Configuring Integrated User Firewall Identification Management for a Tenant System | **636**
- Example: Configure Integrated User Firewall in Customized Model for Tenant System | **646**

The firewall authentication feature is introduced for tenant systems in Junos OS Release 18.3R1 on the Juniper SRX Series devices to enable you to restrict or permit users individually or in groups. The authentication requests are initiated based on destination addresses defined in the policies.

Understanding Tenant System Firewall Authentication

A firewall user is a network user who must provide a username and password for authentication when initiating a connection across the firewall.

Firewall authentication is a policy-based authentication method, which requires user to initiate an authentication request through HTTP, FTP or Telnet traffic.

Junos OS enables administrators to restrict and permit firewall users to access protected resources behind a firewall based on their source IP address and other credentials.

The primary administrator configures the following:

- maximum and reserved number of firewall authentication sessions in the tenant system.
- access profile using the profile configuration command at the `[edit access]` hierarchy which is available to all the tenant systems.

Access profiles allows to:

- Storing usernames and passwords of users or point to external authentication servers where such information is stored.
- Including the order of authentication methods, LDAP or RADIUS server options, and session options.
- Associating with a security policy in the tenant system.

After defining the firewall users, create a policy that requires the users to authenticate through one of the authentication modes defined in the Table 35 on page 618.

Table 35: Firewall Authentication Options

Authentication Options	Description	Supported Protocols	Supported Backend
Web Authentication	Users use HTTP to connect to an IP address on the device that is enabled for Web authentication and are prompted for the username and password. Subsequent traffic from the user or host to the protected resource is allowed or denied based on the result of this authentication.	HTTP HTTPS	Local LDAP RADIUS SecurId
Pass-through	Inline authentication with a host or a user from one zone tries to access resources on another zone. The device uses the supported protocols to collect username and password information, and subsequent traffic from the user or host is allowed or denied based on the result of this authentication.	HTTP HTTPS TELNET FTP	Local LDAP RADIUS SecurId

Table 35: Firewall Authentication Options *(Continued)*

Authentication Options	Description	Supported Protocols	Supported Backend
Web Redirect	Automatically redirect client to WebAuth page for authentication (http or https)	HTTP HTTPS	Local LDAP RADIUS SecurId
Integrated User Firewall	SRX Series devices uses WMI client (WMIC) requests to the AD to get IP address-to-user mapping information in Security event logs.	none	Active Directory
User-Firewall	Same as pass-through but user information is passed to USERID process to go in Auth Table	HTTP HTTPS	Local LDAP RADIUS SecurId

The tenant system administrator configures the following properties for firewall authentication in the tenant system:

- Security policy that specifies firewall authentication for matching traffic. Firewall authentication is specified with the firewall-authentication configuration statement at the `[edit security policies from-zone` *`zone-name`* `to-zone` *`zone-name`* `policy` *`policy-name`* `then permit]` hierarchy level. In an access profile, users or user groups can be allowed access by the policy can optionally be specified with the client-match configuration statement. If no users or user groups are specified, any user who is successfully authenticated is allowed access.

- The type of authentication (pass-through or Web authentication), default access profile, and success banner for the FTP, Telnet, or HTTP session. These properties are configured with the firewall-authentication configuration statement at the `[edit access]` hierarchy.

  Host inbound traffic. Protocols, services, or both are allowed to access the tenant system. The types of traffic are configured with the `host-inbound-traffic`configuration statement at the `[edit security zones security-zone` *`zone-name`*`]` or `[edit security zones security-zone` *`zone-name`* `interfaces` *`interface-name`*`]` hierarchy.

Configuring Firewall Authentication for a Tenant System

IN THIS SECTION

- Requirements | 620
- Overview | 620
- Configuration | 622
- Verification | 633

This example shows how to send different firewall authentication traffic from the client to server across one tenant system using the three authentication modes pass-through, pass-through with web-redirect, and web authentication.

Requirements

This example uses the following hardware and software components:

- an SRX4100 device
- Junos OS Release 18.3R1 and later
- Telnet or HTTP
- External authentication servers are RADIUS, LDAP, and SecurID

Ensure to have the following configured to send firewall authentication traffic from client to server:

- Configure security zones for a tenant system
- Configure interfaces created by the primary administrator

Overview

IN THIS SECTION

- Topology | 622

When a firewall user attempts to initiate a Telnet, HTTP, or HTTPS session to access a resource in another zone, the SRX Series firewall acts a proxy to authenticate the firewall users before allowing the users to access the Telnet, HTTP, or HTTPS servers behind the firewall.

In this example, you can configure a tenant system and bind the security policy to it. When the traffic from is sent from client to server as referred in Figure 18 on page 622, the users are authenticated based on the authentication process defined in the security policy.

NOTE: The primary administrator is responsible for creating tenants and assigning the system resources such as routing-instances, interfaces in routing-instances and security-profile to tenant system.

Table 36: Firewall Configuration for the Tenant System

Feature	Name	Description
security-profile	tn1_pf	Name of the security profile. This profile specifies the resources to allocate to a tenant system to which the security profile is bound.
interfaces	xe-0/0/1 xe-0/0/2	Name of the interfaces. The interfaces provide traffic connectivity.
access profile	local_pf radius_pf securid_pf	Name of the access profiles. These profiles are used to define the users and passwords and to obtain authorization information about the user's access right.
SSL termination profile	fwauthhttpspf	Name of the profile. This profile is used for SSL termination services.
routing-instances	vr1	Instance type as virtual routing instance.
security policies	p7	Name of the policy. This policy is used to configure pass-through firewall-authentication using fwauthhttpspf SSL termination profile.
	p1	Name of the policy. This policy is used to configure pass-through firewall-authentication using local_pf access profile.

Table 36: Firewall Configuration for the Tenant System *(Continued)*

Feature	Name	Description
	p4	Name of the policy. This policy is used to configure pass-through web-redirect firewall-authentication using radius_pf.
	p3	Name of the policy. This policy is used to configure web-authentication firewall-authentication.

Topology

Figure 18 on page 622 shows the topology used in this configuration example. The tenant shown in this topology is an SRX Series device partitioned to multiple tenants. The external servers supported are RADIUS, LDAP, and SecurID. The communication from the client to the tenant happens over xe-0/0/1 interface and from the tenant to the server happens over xe-0/0/2 interface.

Figure 18: Topology for Tenant System

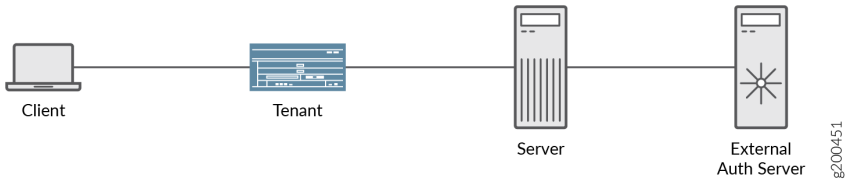

Configuration

IN THIS SECTION

- CLI Quick Configuration | 623
- Configuring access profiles and firewall authentication | 625
- Results | 629

CLI Quick Configuration

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

```
set system security-profile tn1_pf policy maximum 500
set system security-profile tn1_pf policy reserved 100
set system security-profile tn1_pf zone maximum 50
set system security-profile tn1_pf zone reserved 10
set tenants tn1 security-profile tn1_pf
set services ssl termination profile fwauthhttpspf server-certificate device
set interfaces xe-0/0/1 unit 0 family inet address 192.0.2.0/24
set interfaces xe-0/0/1 unit 0 family inet address 192.0.2.254/16 web-authentication http
set interfaces xe-0/0/2 unit 0 family inet address 198.51.100.0/24 web-authentication http
set access profile local_pf client test firewall-user password "$ABC123"
set access profile local_pf client test1 client-group local-group1
set access profile local_pf client test1 client-group local-group2
set access profile local_pf client test1 firewall-user password "$BCD678"
set access profile local_pf client test2 client-group local-group2
set access profile local_pf client test2 firewall-user password "$DEF234"
set access profile local_pf client test3 client-group local-group3
set access profile local_pf client test3 firewall-user password "$DBC123"
set access profile local_pf client test4 client-group local-group4
set access profile local_pf client test4 firewall-user password "$FAB123"
set access profile radius_pf authentication-order radius
set access profile radius_pf radius-server 203.0.113.1 secret "$AFD123"
set access profile securid_pf authentication-order securid
set tenants tn1 routing-instances vr1 instance-type virtual-router
set tenants tn1 routing-instances vr1 interface xe-0/0/1.0
set tenants tn1 routing-instances vr1 interface xe-0/0/2.0
set tenants tn1 security policies from-zone tn1_trust to-zone tn1_untrust policy p1 match source-
address any
set tenants tn1 security policies from-zone tn1_trust to-zone tn1_untrust policy p1 match
destination-address any
set tenants tn1 security policies from-zone tn1_trust to-zone tn1_untrust policy p1 match
application junos-telnet
set tenants tn1 security policies from-zone tn1_trust to-zone tn1_untrust policy p1 then permit
firewall-authentication pass-through access-profile local_pf
set tenants tn1 security policies from-zone tn1_trust to-zone tn1_untrust policy p7 match source-
address any
set tenants tn1 security policies from-zone tn1_trust to-zone tn1_untrust policy p7 match
```

```
destination-address any
set tenants tn1 security policies from-zone tn1_trust to-zone tn1_untrust policy p7 match
application any
set tenants tn1 security policies from-zone tn1_trust to-zone tn1_untrust policy p7 then permit
firewall-authentication pass-through access-profile local_pf
set tenants tn1 security policies from-zone tn1_trust to-zone tn1_untrust policy p7 then permit
firewall-authentication pass-through ssl-termination-profile fwauthhttpspf
set tenants tn1 security policies from-zone tn1_trust to-zone tn1_untrust policy p4 match source-
address any
set tenants tn1 security policies from-zone tn1_trust to-zone tn1_untrust policy p4 match
destination-address any
set tenants tn1 security policies from-zone tn1_trust to-zone tn1_untrust policy p4 match
application junos-http
set tenants tn1 security policies from-zone tn1_trust to-zone tn1_untrust policy p4 then permit
firewall-authentication pass-through access-profile radius_pf
set tenants tn1 security policies from-zone tn1_trust to-zone tn1_untrust policy p4 then permit
firewall-authentication pass-through web-redirect
set tenants tn1 security policies from-zone tn1_trust to-zone tn1_untrust policy p3 match source-
address any
set tenants tn1 security policies from-zone tn1_trust to-zone tn1_untrust policy p3 match
destination-address any
set tenants tn1 security policies from-zone tn1_trust to-zone tn1_untrust policy p3 match
application junos-http
set tenants tn1 security policies from-zone tn1_trust to-zone tn1_untrust policy p3 then permit
firewall-authentication web-authentication
set tenants tn1 security policies policy-rematch
set tenants tn1 security zones security-zone tn1_trust interfaces xe-0/0/1.0 host-inbound-
traffic system-services all
set tenants tn1 security zones security-zone tn1_trust interfaces xe-0/0/1.0 host-inbound-
traffic protocols all
set tenants tn1 security zones security-zone tn1_untrust interfaces xe-0/0/2.0 host-inbound-
traffic system-services all
set tenants tn1 security zones security-zone tn1_untrust interfaces xe-0/0/2.0 host-inbound-
traffic protocols all
set tenants tn1 access firewall-authentication pass-through default-profile local_pf
set tenants tn1 access firewall-authentication pass-through telnet banner login
****tenant1_telnet_login_banner
set tenants tn1 access firewall-authentication pass-through telnet banner success
****tenant1_telnet_success_banner
set tenants tn1 access firewall-authentication pass-through telnet banner fail
****tenant1_telnet_fail_banner
set tenants tn1 access firewall-authentication web-authentication default-profile securid_pf
```

```
set tenants tn1 access firewall-authentication web-authentication banner success
****tenant1_webauth_success_banner
```

Configuring access profiles and firewall authentication

Step-by-Step Procedure

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

1. Configure a security profile tn1_pf and bind it to the tenant system.

```
[edit system security-profile]
user@host# set tn1_pf policy maximum 500
user@host# set tn1_pf policy reserved 100
user@host# set tn1_pf zone maximum 50
user@host# set tn1_pf zone reserved 10
```

2. Create a tenant system tn1 and bind the security profile tn1_pf to the tenant system.

```
[edit tenants]
user@host# set tn1 security-profile tn1_pf
```

3. Define the access profile used for SSL termination services for HTTPS traffic to trigger pass-through authentication.

```
[edit services]
user@host# set ssl termination profile fwauthhttpspf server-certificate device
```

4. Configure interfaces and assign IP addresses. Enable web authentication at xe-0/0/1 interface.

```
[edit interfaces]
user@host# set interfaces xe-0/0/1 unit 0 family inet address 192.0.2.0/24
user@host# set interfaces xe-0/0/1 unit 0 family inet address 192.0.2.254/24 web-
authentication http
user@host# set interfaces xe-0/0/2 unit 0 family inet address 198.51.100.0/24 web-
authentication http
```

5. Configure routing instances and add interfaces to it.

```
[edit tenants tn1 routing-instances]
user@host# set vr1 instance-type virtual-router
user@host# set vr1 interface xe-0/0/1.0
user@host# set vr1 interface xe-0/0/2.0
```

Step-by-Step Procedure

The primary administrator is responsible for configuring access profiles in the tenant system. To configure access profiles:

1. Create the access profiles to be used for firewall authentication. Access profiles defines clients as firewall users and the passwords that provide them access for firewall authentication. When unauthenticated traffic is permitted for firewall authentication, the user is authenticated based on the access profile configured in this command.

```
[edit access profile]
user@host# set local_pf client test firewall-user password "$ABC123"
user@host# set local_pf client test1 client-group local-group1
user@host# set local_pf client test1 client-group local-group2
user@host# set local_pf client test1 firewall-user password "$BCD678"
user@host# set local_pf client test2 client-group local-group2
user@host# set local_pf client test2 firewall-user password "$DEF234"
user@host# set local_pf client test3 client-group local-group3
user@host# set local_pf client test3 firewall-user password "$DBC123"
user@host# set local_pf client test4 client-group local-group4
user@host# set local_pf client test4 firewall-user password "$FAB123"
```

2. Create an access profile to configure the RADIUS server.

```
[edit access profile]
user@host# set radius_pf authentication-order radius
user@host# set radius_pf radius-server 203.0.113.1 secret "$AFD123"
```

3. Create an access profile to configure SecurID as the server to be used for external authentication.

```
[edit access profile]
user@host# set securid_pf authentication-order securid
```

Step-by-Step Procedure

Configure different security policies that permit HTTP, HTTPS, and Telnet traffic between zones using pass-through (direct and web-redirect) and web authentication modes in a tenant system.

1. Configure policy p1 for pass-through authentication for Telnet traffic.

```
[edit tenants tn1 security policies]
user@host# set from-zone tn1_trust to-zone tn1_untrust policy p1 match source-address any
user@host# set from-zone tn1_trust to-zone tn1_untrust policy p1 match destination-address any
user@host# set from-zone tn1_trust to-zone tn1_untrust policy p1 match application junos-
telnet
user@host# set from-zone tn1_trust to-zone tn1_untrust policy p1 then permit firewall-
authentication pass-through access-profile local_pf
```

2. Configure policy p7 for pass-through authentication for HTTPS traffic.

```
[edit tenants tn1 security policies]
user@host# set from-zone tn1_trust to-zone tn1_untrust policy p7 match source-address any
user@host# set from-zone tn1_trust to-zone tn1_untrust policy p7 match destination-address any
user@host# set from-zone tn1_trust to-zone tn1_untrust policy p7 match application junos-https
user@host# set from-zone tn1_trust to-zone tn1_untrust policy p7 then permit firewall-
authentication pass-through access-profile local_pf
user@host# set from-zone tn1_trust to-zone tn1_untrust policy p7 then permit firewall-
authentication pass-through ssl-termination-profile fwauthhttpspf
```

3. Configure policy p4 for pass through authentication using web-redirect for HTTP traffic.

```
[edit tenants tn1 security policies]
user@host# set from-zone tn1_trust to-zone tn1_untrust policy p4 match source-address
ipv6_addr1
user@host# set from-zone tn1_trust to-zone tn1_untrust policy p4 match destination-address any
user@host# set from-zone tn1_trust to-zone tn1_untrust policy p4 match application junos-http
user@host# set from-zone tn1_trust to-zone tn1_untrust policy p4 then permit firewall-
```

```
authentication pass-through access-profile radius_pf
user@host# set from-zone tn1_trust to-zone tn1_untrust policy p4 then permit firewall-
authentication pass-through web-redirect
```

4. Configure policy p3 for web authentication for HTTP traffic.

```
[edit tenants tn1 security policies]
user@host# set from-zone tn1_trust to-zone tn1_untrust policy p3 match source-address any
user@host# set from-zone tn1_trust to-zone tn1_untrust policy p3 match destination-address any
user@host# set from-zone tn1_trust to-zone tn1_untrust policy p3 match application junos-http
user@host# set from-zone tn1_trust to-zone tn1_untrust policy p3 then permit firewall-
authentication web-authentication
user@host# set policy-rematch
```

5. Configure zones and assign interfaces to each zone in a tenant system.

```
[edit tenants tn1 security zones]
user@host# set security-zone tn1_trust interfaces xe-0/0/1.0 host-inbound-traffic system-
services all
user@host# set security-zone tn1_trust interfaces xe-0/0/1.0 host-inbound-traffic protocols
all
user@host# set security-zone tn1_untrust interfaces xe-0/0/2.0 host-inbound-traffic system-
services all
user@host# set security-zone tn1_untrust interfaces xe-0/0/2.0 host-inbound-traffic protocols
all
```

6. Define a success banner for Telnet sessions. Configure firewall authentication pass-through and web authentication banner for applications in a tenant system.

```
[edit tenants tn1 access firewall-authentication]
user@host# set pass-through default-profile local_pf
user@host# set pass-through telnet banner login ****tenant1_telnet_login_banner
user@host# set pass-through telnet banner success ****tenant1_telnet_success_banner
user@host# set pass-through telnet banner fail ****tenant1_telnet_fail_banner
user@host# set web-authentication default-profile securid_pf
user@host# set web-authentication banner success ****tenant1_webauth_success_banner
```

Results

From configuration mode, confirm your configuration by entering the `show system security-profile`, `show interfaces`, `show access`, `show tenants`, and `show services ssl termination` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show interfaces
xe-0/0/1 {
    unit 0 {
        family inet {
            address 192.0.2.0/24;
            address 192.0.2.254/24 {
                web-authentication {
                    http;
                    https;
                }
            }
        }
    }
}
xe-0/0/2 {
    unit 0 {
        family inet {
            address 198.51.100.0/24;
        }
    }
}
```

```
user@host#show services ssl termination
profile fwauthhttpspf {
    server-certificate device;
}
```

```
user@host#show access
profile local_pf {
    client test {
        firewall-user {
            password "$ABC123"; ## SECRET-DATA
        }
```

```
    }
    client test1 {
        client-group [ local-group1 local-group2 ];
        firewall-user {
            password "$BCD678"; ## SECRET-DATA
        }
    }
    client test2 {
        client-group local-group2;
        firewall-user {
            password "$DEF234"; ## SECRET-DATA
        }
    }
    client test3 {
        client-group local-group3;
        firewall-user {
            password "$DBC123"; ## SECRET-DATA
        }
    }
    client test4 {
        client-group local-group4;
        firewall-user {
            password "$FAB123"; ## SECRET-DATA
        }
    }
    session-options {
        client-session-timeout 3;
    }
}
profile radius_pf {
    authentication-order radius;
    session-options {
        client-session-timeout 3;
    }
    radius-server {
        203.0.113.1 secret "$AFD123"; ## SECRET-DATA
    }
}
```

```
user@host#  show system security-profile
tn1_pf {
```

```
    policy {
        maximum 500;
        reserved 100;
    }
    zone {
        maximum 50;
        reserved 10;
    }
}
```

```
user@host#  show tenants
tn1 {
    routing-instances {
        vr1 {
            instance-type virtual-router;
            interface xe-0/0/1.0;
            interface xe-0/0/2.0;
        }
    }
    security-profile {
        tn1_pf;
    }
    security {
        policies {
            from-zone tn1_trust to-zone tn1_untrust {
                policy p2 {
                    match {
                        source-address any;
                        destination-address any;
                        application any;
                    }
                    then {
                        permit {
                            firewall-authentication {
                                pass-through {
                                    access-profile ldap_pf;
                                }
                            }
                        }
                    }
                }
```

```
            }
        }
        zones {
            security-zone tn1_trust {
                interfaces {
                    xe-0/0/1.0 {
                        host-inbound-traffic {
                            system-services {
                                all;
                            }
                            protocols {
                                all;
                            }
                        }
                    }
                }
            }
            security-zone tn1_untrust {
                interfaces {
                    xe-0/0/2.0 {
                        host-inbound-traffic {
                            system-services {
                                all;
                            }
                            protocols {
                                all;
                            }
                        }
                    }
                }
            }
        }
    }
    access {
        firewall-authentication {
            pass-through {
                default-profile local_pf;
                telnet {
                    banner {
                        login ****tenant1_telnet_login_banner;
                        success ****tenant1_telnet_success_banner;
                        fail ****tenant1_telnet_fail_banner;
                    }
```

```
                }
            }
            web-authentication {
                default-profile radius_pf;
                banner {
                    success ****tenant1_webauth_success_banner;
                }
            }
        }
    }
}
```

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

Verification

IN THIS SECTION

- Verifying Firewall User Authentication and Monitoring Users and IP Addresses in the Authentication Table | 633

Verifying Firewall User Authentication and Monitoring Users and IP Addresses in the Authentication Table

Purpose

The administrator for tenant system can use the `show security firewall-authentication users` or `show security firewall-authentication history commands` to view the information about firewall users and history for the tenant system. The administrator for the tenant system can use the same commands to view information for all tenant systems.

Action

From operational mode, enter the following show commands:

```
user@host> show security firewall-authentication history tenant tn1 identifier 10
Username: test
Source IP: 10.12.12.10Authentication state: Success
Authentication method: Pass-through using HTTP
```

```
Access start date: 2018-05-31
Access start time: 17:07:38
Duration of user access: 0:10:01
Lsys: root-logical-system
Tenant: tn1
Source zone: trust-tn1
Destination zone: untrust-tn1
Access profile: test
Bytes sent by this user: 380
Bytes received by this user: 0
user@host> show security firewall-authentication history tenant tn1
History of firewall authentication data:
  Authentications: 2
          Id Source Ip                        Date       Time     Duration   Status   User
           1 203.0.113.10                      2018-05-27 09:33:05 0:01:44    Success  test
           2 203.0.113.10                      2018-05-27 10:01:09 0:10:02    Success  test

user@host> show security firewall-authentication users tenant tn1
Firewall authentication data:
  Total users in table: 1
          Id Source Ip                        Src zone Dst zone Profile    Age Status
User
           2 203.0.113.10                       N/A      N/A      test         1 Success
test
```

Meaning

The output displays the authenticated firewall users and the firewall authentication history of the users for the tenant system

SEE ALSO

firewall-authentication (tenants) | **785**

show security firewall-authentication history | **1011**

show security firewall-authentication users | **1015**

Understanding Integrated User Firewall Support in a Tenant System

IN THIS SECTION

- Limitation of Using User Firewall Authentication in Tenant Systems | 636
- Limitation of using User Firewall Authentication in customized model on Tenant Systems | 636

Tenant system supports the user firewall authentication in shared and active mode.

Starting in Junos OS Release 19.1R1, user firewall authentication is supported on tenant systems using a shared model. In this model, the primary logical system shares the user firewall configuration and authentication entries with the tenant system. The primary logical system shares the authentication data with the tenant system, which is collected from the Local authentication, Active Directory (AD) authentication, firewall authenticationft, Juniper Identity Management Service (JIMS), and ClearPass authentication.

In the shared model, user firewall related configuration is configured under the primary logical system, such as authentication source, authentication source priority, authentication entries timeout, and IP query or individual query and so on. The user firewall provides user information service for an application on the SRX Series device, such as policy and logging. Traffic from a tenant system queries the authentication tables from the primary logical system.

The authentication tables are managed by a primary logical system. The tenant systems share the authentication tables. Traffic from the primary logical system and the tenant systems query the same authentication table. Tenant systems enable the use of the source-identity in security policy.

For example, if the primary logical system is configured with **employee** and the tenant system is configured with the source-identity **manager**, then the reference group of this authentication entry includes **employee** and **manager**. This reference group contains the same authentication entries from primary logical system and tenant system.

Starting in Junos OS Release 19.3R1, support for user firewall authentication is enhanced by using a customized model through integrated JIMS with active mode. In this model, the tenant system extracts the authentication entries from the root level. The primary logical system is configured to the JIMS server based on the logical system and tenant system name. In active mode the SRX series device actively queries the authenticaton entries received from the JIMS server through HTTPs protocol. To reduce the data exchange, firewall filters are applied.

The user firewall uses the tenant system name as a diffrentiator and is consistent between the JIMS server and SRX series device. The JIMS server sends the diffrentiator which is included in the

authentication entry. The authentication entries are distributed into the root logical system, when the diffrentiator is set as default for the primary logical system.

The user firewall support In-service software upgrade (ISSU) for tenant systems, as user firewall changes the internal database table format from Junos OS Release 19.2R1 onwards. Prior to Junos OS Release 19.2R1, the ISSU is not supported for tenant systems.

Starting in Junos OS Release 20.2R1, logical systems and tenant systems support user firewall authentication with Unified Access Control (UAC).

Limitation of Using User Firewall Authentication in Tenant Systems

Using user firewall authentication on tenant systems has the following limitation:

- The IP addresses under different tenant systems must not overlap. If the address overlap, then the authentication entry is changed when different users log in under different tenant systems.

Limitation of using User Firewall Authentication in customized model on Tenant Systems

Using user firewall authentication in customized model on tenant systems has the following limitation:

- The JIMS server configurations to be configured under the root logical systems.
- The tenant system name should be consistent and unique between the JIMS server and the SRX series device.

SEE ALSO

show services user-identification authentication-table

Example: Configuring Integrated User Firewall Identification Management for a Tenant System

IN THIS SECTION

Requirements | 637
Overview | 637

Configuration | 637
Verification | 644

This example shows how to configure the SRX Series device's advanced query feature for obtaining user identity information from the Juniper Identity Management Service (JIMS) and the security policy to match the source identity for a tenant system. In the primary logical system, user firewall is configured with JIMS, and then the primary logical system manages all of authentication entries coming from JIMS. In this example, the primary logical systems shares the authentication entries with the tenant systems.

Requirements

This example uses the following hardware and software components:

- SRX1500 devices operating in chassis clustering
- JIMS server
- Junos OS Release 19.1 R1

Overview

In this example, you can configure JIMS with HTTPs connection on port 443 and primary server with IPv4 address on the primary logical system, policy p1 with source-identity "group1" of dc0 domain on tenant system TN1, policy p1 with source-identity "group1" of dc0 domain on tenant system TN2, and send traffic from and through tenant system TN1 to tenant system TN2. You can view the authentication entries on primary logical system and tenant systems (TN1 and TN2) even after rebooting the primary node.

Configuration

IN THIS SECTION

CLI Quick Configuration | 638
Configuring user firewall identification management | 639
Results | 643

CLI Quick Configuration

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

```
set tenants TN1 security policies from-zone TN1_trust to-zone TN1_trust policy TN1_policy1 match
source-address any
set tenants TN1 security policies from-zone TN1_trust to-zone TN1_trust policy TN1_policy1 match
destination-address any
set tenants TN1 security policies from-zone TN1_trust to-zone TN1_trust policy TN1_policy1 match
application any
set tenants TN1 security policies from-zone TN1_trust to-zone TN1_trust policy TN1_policy1 match
source-identity "example.com\group1"
set tenants TN1 security policies from-zone TN1_trust to-zone TN1_trust policy TN1_policy1 then
permit
set tenants TN1 security policies from-zone TN1_trust to-zone TN1_untrust policy TN1_policy2
match source-address any
set tenants TN1 security policies from-zone TN1_trust to-zone TN1_untrust policy TN1_policy2
match destination-address any
set tenants TN1 security policies from-zone TN1_trust to-zone TN1_untrust policy TN1_policy2
match application any
set tenants TN1 security policies from-zone TN1_trust to-zone TN1_untrust policy TN1_policy2
then permit
set tenants TN1 security policies from-zone TN1_untrust to-zone TN1_trust policy TN1_policy3
match source-address any
set tenants TN1 security policies from-zone TN1_untrust to-zone TN1_trust policy TN1_policy3
match destination-address any
set tenants TN1 security policies from-zone TN1_untrust to-zone TN1_trust policy TN1_policy3
match application any
set tenants TN1 security policies from-zone TN1_untrust to-zone TN1_trust policy TN1_policy3
then permit
set tenants TN1 security policies policy-rematch
set tenants TN2 security policies from-zone TN2_untrust to-zone TN2_untrust policy TN2_policy1
match source-address any
set tenants TN2 security policies from-zone TN2_untrust to-zone TN2_untrust policy TN2_policy1
match destination-address any
set tenants TN2 security policies from-zone TN2_untrust to-zone TN2_untrust policy TN2_policy1
match application any
set tenants TN2 security policies from-zone TN2_untrust to-zone TN2_untrust policy TN2_policy1
match source-identity "example.com\group2"
set tenants TN2 security policies from-zone TN2_untrust to-zone TN2_untrust policy TN2_policy1
```

```
then permit
set tenants TN2 security policies policy-rematch
set services user-identification identity-management connection connect-method https
set services user-identification identity-management connection port 443
set services user-identification identity-management connection primary address 192.0.2.5
set services user-identification identity-management connection primary client-id otest
set services user-identification identity-management connection primary client-secret "$ABC123"
set security policies from-zone root_trust to-zone root_trust policy root_policy1 match source-
address any
set security policies from-zone root_trust to-zone root_trust policy root_policy1 match
destination-address any
set security policies from-zone root_trust to-zone root_trust policy root_policy1 match
application any
set security policies from-zone root_trust to-zone root_trust policy root_policy1 then permit
set security policies policy-rematch
set security zones security-zone root_trust interfaces reth1.0 host-inbound-traffic system-
services all
set security zones security-zone root_trust interfaces reth1.0 host-inbound-traffic protocols all
set security zones security-zone root_trust interfaces lt-0/0/0.1 host-inbound-traffic system-
services all
set security zones security-zone root_trust interfaces lt-0/0/0.1 host-inbound-traffic protocols
all
set firewall family inet filter impair-ldap term allow_all then accept
```

Configuring user firewall identification management

Step-by-Step Procedure

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

To configure user firewall identification management:

1. Log in to the primary logical system as the primary administrator and enter configuration mode.

```
user@host> configure
user@host#
```

2. Create tenant systems.

```
[edit tenants]
user@host#set TN1
user@host#set TN2
```

3. Configure a security policy TN1_policy1 with source-identity group1 on the tenant system TN1 that permits traffic from TN1_trust to TN1_trust.

```
[edit security policies]
user@host#set from-zone TN1_trust to-zone TN1_trust policy TN1_policy1 match source-address
any
user@host#set from-zone TN1_trust to-zone TN1_trust policy TN1_policy1 match destination-
address any
user@host#set from-zone TN1_trust to-zone TN1_trust policy TN1_policy1 match application any
user@host#set from-zone TN1_trust to-zone TN1_trust policy TN1_policy1 match source-
identity "example.com\group1"
user@host#set from-zone TN1_trust to-zone TN1_trust policy TN1_policy1 then permit
```

4. Configure a security policy TN1_policy2 that permits traffic from TN1_trust to TN1_untrust.

```
[edit security policies]
user@host#set from-zone TN1_trust to-zone TN1_untrust policy TN1_policy2 match source-
address any
user@host#set from-zone TN1_trust to-zone TN1_untrust policy TN1_policy2 match destination-
address any
user@host#set from-zone TN1_trust to-zone TN1_untrust policy TN1_policy2 match application
any
user@host#set from-zone TN1_trust to-zone TN1_untrust policy TN1_policy2 then permit
```

5. Configure a security policy TN1_policy3 that permits traffic from TN1_untrust to TN1_trust.

```
[edit security policies]
user@host#set from-zone TN1_untrust to-zone TN1_trust policy TN1_policy3 match source-
address any
user@host#set from-zone TN1_untrust to-zone TN1_trust policy TN1_policy3 match destination-
address any
user@host#set from-zone TN1_untrust to-zone TN1_trust policy TN1_policy3 match application
any
```

```
user@host#set from-zone TN1_untrust to-zone TN1_trust policy TN1_policy3 then permit
user@host#set policy-rematch
```

6. Configure security zone and assign interfaces to each zone.

```
[edit security zones]
user@host#set security-zone TN1_trust interfaces reth2.0 host-inbound-traffic system-
services all
user@host#set security-zone TN1_trust interfaces reth2.0 host-inbound-traffic protocols all
user@host#set security-zone TN1_trust interfaces lt-0/0/0.11 host-inbound-traffic system-
services all
user@host#set security-zone TN1_trust interfaces lt-0/0/0.11 host-inbound-traffic protocols
all
user@host#set security-zone TN1_untrust interfaces reth3.0 host-inbound-traffic system-
services all
user@host#set security-zone TN1_untrust interfaces reth3.0 host-inbound-traffic protocols
all
```

7. Configure a security policy TN2_policy1 with source-identity group1 that permits traffic from TN2_untrust to TN2_untrust on TN2.

```
[edit security policies]
user@host#set from-zone TN2_untrust to-zone TN2_untrust policy TN2_policy1 match source-
address any
user@host#set from-zone TN2_untrust to-zone TN2_untrust policy TN2_policy1 match
destination-address any
user@host#set from-zone TN2_untrust to-zone TN2_untrust policy TN2_policy1 match
application any
user@host#set from-zone TN2_untrust to-zone TN2_untrust policy TN2_policy1 match source-
identity "example.com\group2"
user@host#set from-zone TN2_untrust to-zone TN2_untrust policy TN2_policy1 then permit
user@host#set policy-rematch
```

8. Configure security zones and assign interfaces to each zone on TN2.

```
[edit security zones]
user@host#set security-zone TN2_untrust interfaces reth4.0 host-inbound-traffic system-
services all
user@host#set security-zone TN2_untrust interfaces reth4.0 host-inbound-traffic protocols
all
```

```
user@host#set security-zone TN2_untrust interfaces lt-0/0/0.21 host-inbound-traffic system-
services all
user@host#set security-zone TN2_untrust interfaces lt-0/0/0.21 host-inbound-traffic
protocols all
```

9. Configure JIMS as the authentication source for advanced query requests with the primary address. The SRX Series device requires this information to contact the server.

```
[edit services user-identification identity-management]
user@host#set connection port 443
user@host#set connection connect-method https
user@host#set connection primary address 192.0.2.5
user@host#set connection primary client-id otest
user@host#set connection primary client-secret test
user@host#set authentication-entry-timeout 0
```

10. Configure security policies and zones on the primary logical system.

```
[edit security policies]
user@host#set from-zone root_trust to-zone root_trust policy root_policy1 match source-
address any
user@host#set from-zone root_trust to-zone root_trust policy root_policy1 match destination-
address any
user@host#set from-zone root_trust to-zone root_trust policy root_policy1 match application
any
user@host#set from-zone root_trust to-zone root_trust policy root_policy1 then permit
user@host#set policy-rematch
```

11. Configure security zones and assign interfaces to each zone on primary logical system.

```
[edit security zones]
user@host#set security-zone root_trust interfaces reth1.0 host-inbound-traffic system-
services all
user@host#set security-zone root_trust interfaces reth1.0 host-inbound-traffic protocols all
user@host#set security-zone root_trust interfaces lt-0/0/0.1 host-inbound-traffic system-
services all
user@host#set security-zone root_trust interfaces lt-0/0/0.1 host-inbound-traffic protocols
all
user@host#set firewall family inet filter impair-ldap term allow_all then accept
```

Results

From configuration mode, confirm your configuration by entering the `show services user-identification identity-management` `show chassis cluster` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show services user-identification identity-management
connection {
    connect-method https;
    port 443;
    primary {
        address 192.0.2.5;
        client-id otest;
        client-secret "$ABC123"; ## SECRET-DATA
    }
}
```

```
user@host# show chassis cluster
reth-count 5;
    control-ports {
    fpc 3 port 0;
    fpc 9 port 0;
}
redundancy-group 0 {
    node 0 priority 200;
    node 1 priority 1;
}
redundancy-group 1 {
    node 0 priority 100;
    node 1 priority 1;
}
redundancy-group 2 {
    node 0 priority 100;
    node 1 priority 1;
}
redundancy-group 3 {
    node 0 priority 100;
    node 1 priority 1;
}
redundancy-group 4 {
    node 0 priority 100;
```

```
    node 1 priority 1;
}
```

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

Verification

IN THIS SECTION

- Verifying chassis cluster status and authentication entries | 644
- Verifying chassis cluster status | 645

To confirm that the configuration is working properly, perform the below tasks:

Verifying chassis cluster status and authentication entries

Purpose

To verify authentication entries in a tenant system.

Action

To verify the configuration is working properly, enter the `show services user-identification authentication-table authentication-source identity-management tenant TN1` command.

```
user@host> show services user-identification authentication-table authentication-source identity-
management tenant TN1
node0:
--------------------------------------------------------------------------
Logical System: root-logical-system
Domain: ad2012.jims.com
Total entries: 3
Source IP       Username       groups(Ref by policy)          state
2001:db8:aaaa:  N/A                                           Valid
2001:db8:aaaa:  administrator                                 Valid
203.0.113.50    administrator                                 Valid
node1:
--------------------------------------------------------------------------
```

```
Logical System: root-logical-system
Domain: ad2012.jims.com
Total entries: 3
Source IP       Username       groups(Ref by policy)         state
2001:db8:aaaa:  N/A                                          Valid
2001:db8:aaaa:  administrator                                Valid
203.0.113.50    administrator                                Valid
```

Meaning

The output displays the authentication entries that are shared from the primary logical system to the tenant system.

Verifying chassis cluster status

Purpose

Verify chassis cluster status after rebooting the primary node.

Action

To verify the configuration is working properly, enter the `show chassis cluster status` command.

```
user@host> show chassis cluster status
Monitor Failure codes:
CS  Cold Sync monitoring        FL  Fabric Connection monitoring
GR  GRES monitoring             HW  Hardware monitoring
IF  Interface monitoring        IP  IP monitoring
LB  Loopback monitoring         MB  Mbuf monitoring
NH  Nexthop monitoring          NP  NPC monitoring
SP  SPU monitoring              SM  Schedule monitoring
CF  Config Sync monitoring      RE  Relinquish monitoring
Cluster ID: 6
Node   Priority Status               Preempt Manual   Monitor-failures
Redundancy group: 0 , Failover count: 0
node0  200      hold                 no      no       None
node1  1        secondary            no      no       None
Redundancy group: 1 , Failover count: 0
node0  0        hold                 no      no       CS
node1  1        secondary            no      no       None
Redundancy group: 2 , Failover count: 0
```

```
node0  0        hold                no      no      CS
node1  1        secondary           no      no      None
Redundancy group: 3 , Failover count: 0
node0  0        hold                no      no      CS
node1  1        secondary           no      no      None
Redundancy group: 4 , Failover count: 0
node0  0        hold                no      no      CS
node1  1        secondary           no      no      None
```

Meaning

The output displays user identification management session existing on TN1 and TN2 after rebooting the primary node.

SEE ALSO

show services user-identification authentication-table

Example: Configure Integrated User Firewall in Customized Model for Tenant System

IN THIS SECTION

- Requirements | 647
- Overview | 647
- Configuration | 647
- Verification | 651

This example shows how to configure the integrated user firewall by using a customized model through the Juniper Identity Management Service (JIMS) server with active mode for a tenant system. The primary logical systems does not share the authentication entries with the tenant systems. The SRX series device queries the authentication entries received from the JIMS server through HTTPs protocol in active mode.

In this example following configurations are performed:

- Active JIMS Server Configuration
- Tenant System IP Query Configuration
- Tenant System Authentication Entry Configuration
- Tenant System Security Policy Configuration

Requirements

This example uses the following hardware and software components:

- JIMS server version 2.0
- Junos OS Release 19.3R1

Before you begin, be sure you have following information:

- The IP address of the JIMS server.
- The port number on the JIMS server for receiving HTTPs requests.
- The client ID from the JIMS server for active query server.
- The client secret from the JIMS server for active query server.

Overview

In this example, you can configure JIMS with HTTPs connection on port 443 and primary server with IPv4 address on the primary logical system, policy p2 with source-identity `group1` on tenant system `TSYS1`.

Configuration

IN THIS SECTION

- CLI Quick Configuration | **648**
- Configuring Integrated User Firewall in Customized Model: | **648**
- Results | **650**

CLI Quick Configuration

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

```
set services user-identification logical-domain-identity-management active query-server jims1
connection connect-method https
set services user-identification logical-domain-identity-management active query-server jims1
connection port 443
set services user-identification logical-domain-identity-management active query-server jims1
connection primary address 192.0.2.5
set services user-identification logical-domain-identity-management active query-server jims1
connection primary client-id otest
set services user-identification logical-domain-identity-management active query-server jims1
connection primary client-secret "$ABC123"
set tenants TSYS1 services user-identification logical-domain-identity-management active ip-
query query-delay-time 30
set tenants TSYS1 services user-identification logical-domain-identity-management active invalid-
authentication-entry-timeout 1
set tenants TSYS1 security policies from-zone untrust to-zone trust policy p2 match source-
address any
set tenants TSYS1 security policies from-zone untrust to-zone trust policy p2 match destination-
address any
set tenants TSYS1 security policies from-zone untrust to-zone trust policy p2 match application
any
set tenants TSYS1 security policies from-zone untrust to-zone trust policy p2 match source-
identity "example.com\group1"
set tenants TSYS1 security policies from-zone untrust to-zone trust policy p2 then permit
```

Configuring Integrated User Firewall in Customized Model:

Step-by-Step Procedure

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

To configure Integrated User Firewall in Customized Model:

1. Configure JIMS as the authentication source for advanced query requests with the primary address. The SRX Series device requires this information to contact the server.

```
user@host# set services user-identification logical-domain-identity-management active query-server jims1 connection connect-method https
user@host# set services user-identification logical-domain-identity-management active query-server jims1 connection port 443
user@host# set services user-identification logical-domain-identity-management active query-server jims1 connection primary address 192.0.2.5
user@host# set services user-identification logical-domain-identity-management active query-server jims1 connection primary client-id otest
user@host# set services user-identification logical-domain-identity-management active query-server jims1 connection primary client-secret "$ABC123"
```

2. Configure the IP query delay time for TSYS1.

```
user@host# set tenants TSYS1 services user-identification logical-domain-identity-management active ip-query query-delay-time 30
```

3. Configure the authentication entry attributes for TSYS1.

```
user@host# set tenants TSYS1 services user-identification logical-domain-identity-management active invalid-authentication-entry-timeout 1
```

4. Configure the security policy p2 that permits traffic from-zone untrust to-zone trust for TSYS1.

```
user@host# set tenants TSYS1 security policies from-zone untrust to-zone trust policy p2 match source-address any
user@host# set tenants TSYS1 security policies from-zone untrust to-zone trust policy p2 match destination-address any
user@host# set tenants TSYS1 security policies from-zone untrust to-zone trust policy p2 match application any
user@host# set tenants TSYS1 security policies from-zone untrust to-zone trust policy p2 match source-identity "example.com\group1"
user@host#set tenants TSYS1 security policies from-zone untrust to-zone trust policy p2 then permit
```

Results

From configuration mode, confirm your configuration by entering the `show services user-identification logical-domain-identity-management` and `show tenants TSYS1` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show services user-identification logical-domain-identity-management
active {
    query-server jims1 {
        connection {
            connect-method https;
            port 443;
            primary {
                address 10.1.1.1;
                client-id otest;
                client-secret "$ABC123"; ## SECRET-DATA
            }
        }
    }
}
```

```
user@host# show tenants TSYS1
security {
    policies {
        from-zone untrust to-zone trust {
            policy p2 {
                match {
                    source-address any;
                    destination-address any;
                    application any;
                    source-identity "example.com\group1";
                }
                then {
                    permit;
                }
            }
        }
    }
}
services {
    user-identification {
```

```
        logical-domain-identity-management {
            active {
                invalid-authentication-entry-timeout 1;
                ip-query {
                    query-delay-time 30;
                }
            }
        }
    }
}
```

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

Verification

IN THIS SECTION

- Verifying the User Identification Identity Management status | **651**
- Verifying the User Identification Identity Management status counters | **652**
- Verifying the User Identification Authentication Table | **653**

To confirm that the configuration is working properly, perform the below tasks:

Verifying the User Identification Identity Management status

Purpose

Verify the user identification status for identity-management as the authentication source.

Action

To verify the configuration is working properly, enter the `show services user-identification logical-domain-identity-management status` command.

```
user@host>show services user-identification logical-domain-identity-management status
node0:
--------------------------------------------------------------------------
      Query server name                :jims1
```

```
      Primary server :
     Address                      : 10.1.1.1
     Port                         : 443
     Connection method            : HTTPS
     Connection status            : Online
     Last received status message : OK (200)
     Access token                 : isdHIbl8BXwxFftMRubGVsELRukYXtW3rtKmHiL
     Token expire time            : 2017-11-27 23:45:22
     Secondary server :
     Address                      : Not configured
```

Meaning

The output displays the statistical data about the advanced user query function batch queries and IP queries, or show status on the Juniper Identity Management Service servers.

Verifying the User Identification Identity Management status counters

Purpose

Verify the user identification counters for identity-management as the authentication source.

Action

To verify the configuration is working properly, enter the `show services user-identification logical-domain-identity-management counters` command.

```
user@host>show services user-identification logical-domain-identity-management counters
node0:
      --------------------------------------------------------------------------
   Query server name                      :jims1
    Primary server :
     Address                           : 10.208.137.208
     Batch query sent number           : 65381
     Batch query total response number : 64930
     Batch query error response number : 38
     Batch query last response time    : 2018-08-14 15:10:52
     IP query sent number              : 10
     IP query total response number    : 10
     IP query error response number    : 0
     IP query last response time       : 2018-08-13 12:41:56
```

```
  Secondary server :
 Address                          : Not configured
```

Meaning

The output displays the statistical data about the advanced user query function batch queries and IP queries, or show counters on the Juniper Identity Management Service servers.

Verifying the User Identification Authentication Table

Purpose

Verify the user identity information authentication table entries for the specified authentication source.

Action

To verify the configuration is working properly, enter the `show services user-identification authentication-table authentication-source all tenant TSYS1` command.

```
user@host>show services user-identification authentication-table authentication-source all tenant TSYS1
node0:
      --------------------------------------------------------------------------
   Tenant System: TSYS1
Domain: ad03.net
Total entries: 4
Source IP       Username       groups(Ref by policy)          state
10.12.0.2        administrator  posture-healthy                Valid
10.12.0.15       administrator  posture-healthy                Valid
2001:db8:3000::5         N/A            posture-healthy                Valid
fe80::342c:302b N/A            posture-healthy                Valid
```

Meaning

The output displays the entire content of the specified authentication source's authentication table, or a specific domain, group, or user based on the user name. Display the identity information for a user based on the IP address of the user's device.

Release History Table

Release	Description
19.3R1	Starting in Junos OS Release 19.3R1, support for user firewall authentication is enhanced by using a customized model through integrated JIMS with active mode.
19.1R1	Starting in Junos OS Release 19.1R1, user firewall authentication is supported on tenant systems using a shared model. In this model, the primary logical system shares the user firewall configuration and authentication entries with the tenant system. The primary logical system shares the authentication data with the tenant system, which is collected from the Local authentication, Active Directory (AD) authentication, firewall authenticationft, Juniper Identity Management Service (JIMS), and ClearPass authentication.

Security Policies for Tenant Systems

IN THIS SECTION

- Understanding Security Policies for Tenant Systems | **654**
- Example: Configuring Security Policies in the Tenant System | **656**
- Configuring Dynamic Address for Tenant Systems | **662**

Security policies can be configured with tenant systems. For more information see the following topics:

Understanding Security Policies for Tenant Systems

IN THIS SECTION

- Application Timeouts | **655**
- Security Policy Allocation | **656**

Security policies enforce rules for what traffic can pass through the firewall and actions that need to take place on the traffic as it passes through the firewall. Through the creation of security policies, the administrator for the tenant system can control the traffic flow from zone to zone by defining the kinds of traffic permitted to pass from sources to destinations. From the perspective of the security policies, traffic enters one security zone and exits through another security zone. By default, the tenant system denies all traffic in all directions, including intra-zone and inter-zone directions.

Starting in Junos OS Release 18.3R1, the security policies feature supported on logical systems is now extended to tenant systems.

Security policies can be configured in the tenant systems. Tenant security policies are configured the same way as logical system security policies and firewall-wide security policies. Any security policies, policy rules, address books, applications and application sets, and schedulers created within a tenant system are only applicable to that tenant system. Only predefined applications and application sets, such as `junos-ftp`, are shared between the tenant systems.

The administrator for the tenant system can configure and view all attributes for security policies in a tenant system.

Starting in Junos OS Release 18.4R1, the tenant system administrator can create dynamic address within a tenant system. A dynamic address entry contains IP addresses and prefixes extracted from external sources. The security policies use the dynamic address in the source-address field or destination-address field. You can view the dynamic-address information including the name, feeds, and properties for tenant systems by using the command `show security dynamic-address`.

A dynamic address entry (DAE) is a group of IP addresses that can be entered manually or imported from external sources within tenant systems. The DAE feature allows feed-based IP objects to be used in security policies to either deny or allow traffic based on either source or destination IP criteria.

NOTE: The maximum number of DAE for a given tenant system equals the system-wide scaling number. Furthermore, the sum of DAE for all the tenant systems must be less than or equal to the system-wide scaling number for DAE. If one tenant system uses maximum number of IP entries, other tenant system will fail to get IP entries into their DAE.

Starting in Junos 18.4R1, the `set security dynamic-address feed-server` command can be configured under the tenant systems.

Application Timeouts

The application timeout value set for an application determines the session timeout. Application timeout behavior is the same for a tenant system as it is at the root level. Although the administrators of the tenant system can use predefined applications in security policies, the administrators cannot modify the timeout value for these predefined applications. Application timeout values are stored in the application entry database and in the corresponding tenant system TCP and UDP port-based timeout tables.

Security Policy Allocation

The primary administrator creates a security profile to allocate the maximum number of policies that can be configured for each tenant system. The administrator of the tenant system is then restricted by the security profile to create no more than the number of policies described in the security profile. The administrator of the tenant system use the `show system security-profile policy` command to view the number of security policies allocated to the tenant system.

```
user@host> show system security-profile policy
```

```
logical-system tenant name   security profile name       usage    reserved     maximum

root-logical-system          Default-Profile                1           0       16000
```

Example: Configuring Security Policies in the Tenant System

IN THIS SECTION

- Requirements | 656
- Overview | 657
- Configuration | 658
- Verification | 660

This example shows how to configure the security policies for the tenant system.

Requirements

Before you begin the configuration:

- Configure zones. See *Example: Configuring Security Zones in the Tenant System*.
- Use the `show system security-profiles policy` command to see the security policy resources allocated to the tenant system.

Overview

In this example, you can configure a security policy for the tenant system. The administrator for the tenant system user can use `[edit tenants tenant-name security policies]` hierarchy level to configure the security policies. This example configures the security policies described in Table 37 on page 657.

Table 37: Security Policies Parameters

Feature	Configuration Parameters
Policy 1	Permit the following traffic: • Policy name: p1 • Tenant name: TSYS1 • From zone: trust • To zone: untrust • Source address: any • Destination address: any • Application: any
Policy 2	Permit the following traffic: • Policy name: p1 • Tenant name: TSYS1 • From zone: untrust • To zone: trust • Source address: any • Destination address: any • Application: any

Configuration

IN THIS SECTION

- Procedure | 658

Procedure

CLI Quick Configuration

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

```
set tenants TSYS1 security policies from-zone trust to-zone untrust policy p1 match source-
address any
set tenants TSYS1 security policies from-zone trust to-zone untrust policy p1 match destination-
address any
set tenants TSYS1 security policies from-zone trust to-zone untrust policy p1 match application
any
set tenants TSYS1 security policies from-zone trust to-zone untrust policy p1 then permit
set tenants TSYS1 security policies from-zone untrust to-zone trust policy p2 match source-
address any
set tenants TSYS1 security policies from-zone untrust to-zone trust policy p2 match destination-
address any
set tenants TSYS1 security policies from-zone untrust to-zone trust policy p2 match application
any
set tenants TSYS1 security policies from-zone untrust to-zone trust policy p2 then permit
```

Step-by-Step Procedure

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

To configure the security policies in the tenant system:

1. Log in to the tenant system and define the tenant system name as TSYS1.

```
[edit]
user@host# set tenants TSYS1
```

2. Create a security policy as p1 that permits traffic from zone trust to zone untrust and configure the match condition.

```
[edit tenants TSYS1 security policies from-zone trust to-zone untrust]
user@host# set policy p1 match source-address any
user@host# set policy p1 match destination-address any
user@host# set policy p1 match application any
user@host# set policy p1 then permit
```

3. Create a security policy as p2 that permits traffic from zone untrust to zone trust and configure the match condition.

```
[edit tenants TSYS1 security policies from-zone untrust to-zone trust]
user@host# set policy p2 match source-address any
user@host# set policy p2 match destination-address any
user@host# set policy p2 match application any
user@host# set policy p2 then permit
```

Results

From configuration mode, confirm your configuration by entering the `show tenants tenant-name security policies` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show tenants TSYS1 security policies
from-zone trust to-zone untrust {
    policy p1 {
        match {
            source-address any;
            destination-address any;
            application any;
        }
        then {
```

```
            permit;
        }
    }
}
from-zone untrust to-zone trust {
    policy p2 {
        match {
            source-address any;
            destination-address any;
            application any;
        }
        then {
            permit;
        }
    }
}
```

Verification

IN THIS SECTION

- Verifying Policy Configuration | 660

Verifying Policy Configuration

Purpose

Verify the information about security policies.

Action

To verify the configuration is working properly, enter the `show security policies detail tenant TSYS1` command from operational mode.

```
user@host> show security policies detail tenant TSYS1
```

```
Default policy: deny-all
Pre ID default policy: permit-all
Policy: p1, action-type: permit, State: enabled, Index: 4, Scope Policy: 0
Policy Type: Configured
Sequence number: 1
From zone: trust, To zone: untrust
Source addresses: any
Destination addresses: any
Application: any
IP protocol: 1, ALG: 0, Inactivity timeout: 60
ICMP Information: type=255, code=0
Application: junos-telnet
IP protocol: tcp, ALG: 0, Inactivity timeout: 1800
Source port range: [0-0]
Destination port range: [23-23]
Application: app_udp
IP protocol: udp, ALG: 0, Inactivity timeout: 1800
Source port range: [0-0]
Destination port range: [5000-5000]
Application: junos-icmp6-all
IP protocol: 58, ALG: 0, Inactivity timeout: 60
ICMP Information: type=255, code=0
Per policy TCP Options: SYN check: No, SEQ check: No, Window scale: No
Session log: at-create, at-close
Policy statistics:
Input  bytes     :                   0                   0 bps
Initial direction:                   0                   0 bps
Reply direction  :                   0                   0 bps
Output bytes     :                   0                   0 bps
Initial direction:                   0                   0 bps
Reply direction  :                   0                   0 bps
Input  packets   :                   0                   0 pps
Initial direction:                   0                   0 bps
```

```
Reply direction   :                0                    0 bps
Output packets    :                0                    0 pps
Initial direction:                 0                    0 bps
Reply direction   :                0                    0 bps
Session rate      :                0                    0 sps
Active sessions   :                0
Session deletions:                 0
Policy lookups    :                0
```

Meaning

The output displays the information about the security policies configured on the tenant system.

Configuring Dynamic Address for Tenant Systems

A dynamic address entry in the tenant system provides dynamic IP address information to security policies. To use dynamic address, you must specify basic information of dynamic address including their names, feeds and properties for a tenant system.

- Read the "Example: Configuring Security Policies in the Tenant System" on page 656 to understand how and where this procedure fits in the overall tenant support for security policy.

To configure the dynamic address in IPv4 networks within a tenant system:

1. Define the tenant system name as TSYS1.

```
[edit]
user@host# set tenants TSYS1
```

2. Create dynamic address within a tenant system.

```
[edit tenants TSYS1]
user@host# set security dynamic-address address-name Ipv4 profile category IPFilter feed fd1
```

3. Confirm your configuration by entering the `show tenants TSYS1 security dynamic-address` command.

```
[edit]
user@host# show tenants TSYS1 security dynamic-address
address-name Ipv4 {
```

```
    profile {
        category GeoIP;
        category IPFilter {
            feed fd1;
            }
        }
    }
}
```

- To configure the security policies in the tenant system:

  1. Define the tenant system name as TSYS1.

     ```
     [edit]
     user@host# set tenants TSYS1
     ```

  2. Create a security policy as p1 that permits traffic from zone trust to zone untrust and configure the match condition.

     ```
     [edit tenants TSYS1 security policies from-zone trust to-zone untrust]
     user@host# set policy p1 match source-address any
     user@host# set policy p1 match destination-address any
     user@host# set policy p1 match application any
     user@host# set policy p1 then permit
     ```

  3. Confirm your configuration by entering the `show tenants tenant-name security policies` command

     ```
     [edit]
     user@host# show tenants TSYS1 security policies
     from-zone trust to-zone untrust {
         policy p1 {
             match {
                 source-address any;
                 destination-address any;
                 application any;
             }
             then {
                 permit;
             }
     ```

```
    }
}
```

Release History Table

Release	Description
18.3R1	Starting in Junos OS Release 18.3R1, the security policies feature supported on logical systems is now extended to tenant systems.

RELATED DOCUMENTATION

Tenant Systems Overview | 530

Screen Options for Tenant Systems

IN THIS SECTION

- Understanding Tenant System Screen Options | 664
- Example: Configuring Screen Options for a Tenant System | 665

Screen options for Tenant Systems on SRX Series devices prevent attacks as , such as IP address sweeps, port scans, denial of service (DOS) attacks, ICMP, UDP, and SYN floods as same as Logical Systems. For more information, see the following topics:

Understanding Tenant System Screen Options

Using screen options, the device secures a zone by inspecting, and then allowing or denying all connection attempts that require crossing an interface bound to that zone. Junos OS applies the firewall policies, which can contain the content filtering and the IDP components to the traffic that passes the screen filters. All screen options that are available on the device are also available in each tenant system.

Starting in Junos OS Release 18.3R1, the screen options that are supported for logical systems are extended to tenant systems.

SEE ALSO

Understanding Screens Options on SRX Series Devices

Example: Configuring Screen Options for a Tenant System

IN THIS SECTION

- Requirements | 665
- Overview | 665
- Configuration | 666
- Verification | 670

This example shows how to configure screen options for a tenant system.

Requirements

Before you begin:

- Understand the tenant system configuration process. See "Tenant System Configuration Overview" on page 537 to understand how this task fits into the overall configuration process.
- Configure the zones for the tenant system. See "Security Zones for Tenant Systems" on page 572to understand how to configure the zones for the tenant systems.

Overview

Using screen options, the security device can protect against the different internal and external attacks for security zones. You can limit the number of concurrent sessions to the same destination IP address in a tenant system. Setting a destination based session limit can ensure that Junos OS allows only an acceptable number of concurrent connection requests—no matter what the source—to reach any one host. When the number of concurrent connection requests to an IP address surpasses the limit, Junos OS blocks further connection attempts to that IP address.

Configuration

IN THIS SECTION

- Procedure | 666

Procedure

CLI Quick Configuration

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

```
set tenants TN1 security screen ids-option jscreen limit-session destination-ip-based 80
set tenants TN1 security screen ids-option jscreen icmp ip-sweep threshold 1000
set tenants TN1 security screen ids-option jscreen icmp fragment
set tenants TN1 security screen ids-option jscreen icmp large
set tenants TN1 security screen ids-option jscreen icmp flood threshold 200
set tenants TN1 security screen ids-option jscreen icmp ping-death
set tenants TN1 security screen ids-option jscreen ip bad-option
set tenants TN1 security screen ids-option jscreen ip stream-option
set tenants TN1 security screen ids-option jscreen ip spoofing
set tenants TN1 security screen ids-option jscreen ip strict-source-route-option
set tenants TN1 security screen ids-option jscreen ip unknown-protocol
set tenants TN1 security screen ids-option jscreen ip tear-drop
set tenants TN1 security screen ids-option jscreen tcp syn-fin
set tenants TN1 security screen ids-option jscreen tcp tcp-no-flag
set tenants TN1 security screen ids-option jscreen tcp syn-frag
set tenants TN1 security screen ids-option jscreen tcp port-scan threshold 1000
set tenants TN1 security screen ids-option jscreen tcp syn-ack-ack-proxy threshold 500
set tenants TN1 security screen ids-option jscreen tcp syn-flood alarm-threshold 500
set tenants TN1 security screen ids-option jscreen tcp syn-flood attack-threshold 500
set tenants TN1 security screen ids-option jscreen tcp syn-flood source-threshold 50
set tenants TN1 security screen ids-option jscreen tcp syn-flood destination-threshold 1000
set tenants TN1 security screen ids-option jscreen tcp syn-flood timeout 10
set tenants TN1 security screen ids-option jscreen tcp land
set tenants TN1 security screen ids-option jscreen tcp winnuke
set tenants TN1 security screen ids-option jscreen tcp tcp-sweep threshold 1000
```

```
set tenants TN1 security screen ids-option jscreen udp flood threshold 500
set tenants TN1 security screen ids-option jscreen udp udp-sweep threshold 1000
set tenants TN1 security zones security-zone untrust screen jscreen
```

Step-by-Step Procedure

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

To configure destination-based session limits in a tenant system:

1. Log in to the tenant system as the administrator and enter configuration mode.

```
user@host:TN1#> configure
user@host:TN1#
```

2. Define the tenant system name as TN1 and configure a screen option for a destination-based session limit.

```
[edit tenants TN1]
user@host:TN1# set security screen ids-option jscreen limit-session destination-ip-based 80
```

3. Configure the ICMP screening options.

```
[edit tenants TN1 security screen ids-option jscreen]
user@host:TN1# set icmp ip-sweep threshold 1000
user@host:TN1# set icmp fragment
user@host:TN1# set icmp large
user@host:TN1# set icmp flood threshold 200
user@host:TN1# set icmp ping-death
```

4. Configure the IP screening options.

```
[edit tenants TN1 security screen ids-option jscreen]
user@host:TN1# set ip bad-option
user@host:TN1# set ip stream-option
user@host:TN1# set ip spoofing
user@host:TN1# set ip strict-source-route-option
```

```
user@host:TN1# set ip unknown-protocol
user@host:TN1# set ip tear-drop
```

5. Configure the TCP screening options.

```
[edit tenants TN1 security screen ids-option jscreen]
user@host:TN1# set tcp syn-fin
user@host:TN1# set tcp tcp-no-flag
user@host:TN1# set tcp syn-frag
user@host:TN1# set tcp port-scan threshold 1000
user@host:TN1# set tcp syn-ack-ack-proxy threshold 500
user@host:TN1# set tcp syn-flood alarm-threshold 500
user@host:TN1# set tcp syn-flood attack-threshold 500
user@host:TN1# set tcp syn-flood source-threshold 50
user@host:TN1# set tcp syn-flood destination-threshold 1000
user@host:TN1# set tcp syn-flood timeout 10
user@host:TN1# set tcp land
user@host:TN1# set tcp winnuke
user@host:TN1# set tcp tcp-sweep threshold 1000
```

6. Configure the UDP screening options.

```
[edit tenants TN1 security screen ids-option jscreen]
user@host:TN1# set udp flood threshold 500
user@host:TN1# set udp udp-sweep threshold 1000
```

7. Attach the IDS profile to the zone.

```
[edit tenants TN1]
user@host:TN1# set security zones security-zone untrust screen jscreen
```

Results

From configuration mode, confirm your configuration by entering the `show tenants TN1 security screen` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show tenants TN1 security screen
ids-option jscreen {
    limit-session {
        destination-ip-based 80;
    }
}
ids-option jscreen {
    icmp {
        ip-sweep threshold 1000;
        fragment;
        large;
        flood threshold 200;
        ping-death;
    }
    ip {
        bad-option;
        stream-option;
        spoofing;
        strict-source-route-option;
        unknown-protocol;
        tear-drop;
    }
    tcp {
        syn-fin;
        tcp-no-flag;
        syn-frag;
        port-scan threshold 1000;
        syn-ack-ack-proxy threshold 500;
        syn-flood {
            alarm-threshold 500;
            destination-threshold 1000;
            timeout 10;
        }
        land;
        winnuke;
        tcp-sweep threshold 1000;
```

```
    }
    udp {
        flood {
            threshold 500;
        }
        udp-sweep threshold 1000;
    }
}
```

Verification

IN THIS SECTION

- Verifying security screen status | 670

To confirm that the configuration is working properly, perform the below task:

Verifying security screen status

Purpose

Verify that the IDS profile for multiple screening options is configured properly:

Action

To verify the configuration is working properly, enter the `show security screen ids-option jscreen tenant TN1` and `show security zone tenant TN1`command from operational mode.

```
user@host> show security screen ids-option jscreen tenant TN1

  Screen object status:

  Name                                    Value
  ICMP flood threshold                    200
  UDP flood threshold                     500
  TCP winnuke                             enabled
  TCP port scan threshold                 1000
  ICMP address sweep threshold            1000
```

```
  TCP sweep threshold                     1000
  UDP sweep threshold                     1000
  IP tear drop                            enabled
  TCP SYN flood attack threshold          500
  TCP SYN flood alarm threshold           500
  TCP SYN flood source threshold          50
  TCP SYN flood destination threshold     1000
  TCP SYN flood timeout                   10
  IP spoofing                             enabled
  ICMP ping of death                      enabled
  TCP land attack                         enabled
  TCP SYN fragment                        enabled
  TCP no flag                             enabled
  IP unknown protocol                     enabled
  IP bad options                          enabled
  IP strict source route option           enabled
  IP stream option                        enabled
  ICMP fragmentation                      enabled
  ICMP large packet                       enabled
  TCP SYN FIN                             enabled
  TCP SYN-ACK-ACK proxy threshold         500
```

```
user@host> show security zone tenant TN1

Security zone: untrust
Send reset for non-SYN session TCP packets: Off
Policy configurable: Yes
Screen: jscreen
Interfaces bound: 0
Interfaces:
```

Meaning

The output displays the screen status in the tenant system.

SEE ALSO

Understanding Tenant System Screen Options | 664

Release History Table

Release	Description
18.3R1	Starting in Junos OS Release 18.3R1, the screen options that are supported for logical systems are extended to tenant systems.

RELATED DOCUMENTATION

Example: Creating Tenant Systems, Tenant System Administrators, and an Interconnect VPLS Switch

NAT for Tenant Systems

IN THIS SECTION

- Understanding Network Address Translation for Tenant systems | 672
- Example: Configuring Network Address Translation for the Tenant Systems | 673

NAT is a method for modifying or translating network address information in packet headers. Either or both source and destination addresses in a packet may be translated. For more information, see the following topics:

Understanding Network Address Translation for Tenant systems

Starting in Junos OS Release 18.3R1, the network address translation including source NAT, destination NAT, and static NAT supported on logical systems is supported on tenant systems.

A tenant system has an administrator (tenant administrator) who can configure source NAT, destination NAT, and static NAT for the tenant systems. The tenant administrator can view the details of the source NAT, destination NAT, and static NAT of the tenant system. The primary administrator can view the statistics or information of the source NAT, destination NAT, and static NAT for any tenant systems.

For the tenant system, the primary administrator can configure the maximum and reserved numbers for the following NAT resources:

- Source NAT pools and destination NAT pools
- IP addresses in the source NAT pools with and without port address translation
- Rules for source, destination, and static NAT
- Prefix list for rule matching
- NAT cone binding
- IP addresses that support port overloading

The reserved numbers allocated guarantees that the specified resource amount is constantly available to the tenant systems. The administrator for tenant systems can use the `show system security-profile` command with a NAT option to view the NAT resources allocated to the tenant system.

SEE ALSO

Understanding Network Address Translation for Tenant systems | 672

Introduction to NAT

Example: Configuring Network Address Translation for the Tenant Systems

IN THIS SECTION

- Requirements | 673
- Overview | 674
- Configuration | 675
- Verification | 679

This example shows how to configure source NAT, destination NAT and static NAT for a given tenant systems.

Requirements

This example uses the following hardware and software components:

- SRX Series device with Junos OS Release 18.3R1 or later. This configuration example is tested for Junos OS Release 18.3R1.
- Create tenant system. See : Example: Creating Tenant Systems, Tenant System Administrators, and an Interconnect VPLS Switch
- Configure network interfaces. See : "Configuring a Routing Instance for a Tenant System" on page 539.

Overview

In this example, first you configure the trust security zone for the private address space and then you configure the untrust security zone for the public address space.

Devices in the untrust zone access a specific host in the trust zone, with the destination IP address 203.0.113.200/24. This example configures the NAT described in Table 1: Tenant System NAT Configuration.

Table 38: Tenant System NAT Configuration

Feature	Name	Configuration Parameters
Static, source and destination NAT rule set	r1	• Rule r1 to match packets from untrust zone with destination address. • Destination IP address in matching packets is translated.
Source pool	pat	Address 192.0.2.1 to 192.0.2.24.
Destination pool	h1	Address 192.168.1.200.
Proxy ARP	arp	Address 192.0.2.1 to 192.0.2.24.
NAT interfaces for traffic direction.		ge-0/0/0 and ge-0/0/1.

Configuration

IN THIS SECTION

- Procedure | 675

Procedure

CLI Quick Configuration

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

```
set tenants tn1 security nat source pool pat address 192.0.2.1 to 192.0.2.24
set tenants tn1 security nat source rule-set from_intf from interface ge-0/0/0.0
set tenants tn1 security nat source rule-set from_intf to interface ge-0/0/1.0
set tenants tn1 security nat source rule-set from_intf rule r1 match source-address 192.0.2.0/24
set tenants tn1 security nat source rule-set from_intf rule r1 match destination-address
203.0.113.200/24
set tenants tn1 security nat source rule-set from_intf rule r1 then source-nat pool pat
set tenants tn1 security nat static rule-set from_zone from zone trust
set tenants tn1 security nat static rule-set from_zone rule r1 match source-address 192.0.2.0/24
set tenants tn1 security nat static rule-set from_zone rule r1 match destination-address
203.0.113.203/24
set tenants tn1 security nat static rule-set from_zone rule r1 then static-nat prefix
192.168.1.203/24
set tenants tn1 security nat destination pool h1 address 192.168.1.200
set tenants tn1 security nat destination rule-set from_zone from zone trust
set tenants tn1 security nat destination rule-set from_zone rule r1 match source-address
192.0.2.0/24
set tenants tn1 security nat destination rule-set from_zone rule r1 match destination-address
203.0.113.202/24
set tenants tn1 security nat destination rule-set from_zone rule r1 then destination-nat pool h1
set tenants tn1 security nat proxy-arp interface ge-0/0/1.0 address 192.0.2.1 to 192.0.2.24
```

Step-by-Step Procedure

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

To configure NAT in the tenant system:

1. Create a security NAT source pool and rule set for the tenant system.

```
[edit tenant tn1 security nat source]
user@host# set tenants tn1 security nat source pool pat address 192.0.2.1 to 192.0.2.24
user@host# set tenants tn1 security nat source rule-set from_intf from interface ge-0/0/0.0
user@host# set tenants tn1 security nat source rule-set from_intf to interface ge-0/0/1.0
user@host# set tenants tn1 security nat source rule-set from_intf rule r1 match source-
address 192.0.2.0/24
user@host# set tenants tn1 security nat source rule-set from_intf rule r1 match destination-
address 203.0.113.200/24
user@host# set tenants tn1 security nat source rule-set from_intf rule r1 then source-nat
pool pat
```

2. Create a security NAT static rule set for the tenant system.

```
[edit tenants tn1 security nat static]
user@host# set tenants tn1 security nat static rule-set from_zone from zone trust
user@host# set tenants tn1 security nat static rule-set from_zone rule r1 match source-
address 192.0.2.0/24
user@host# set tenants tn1 security nat static rule-set from_zone rule r1 match destination-
address 203.0.113.203/24
user@host# set tenants tn1 security nat static rule-set from_zone rule r1 then static-nat
prefix 192.168.1.203/24
```

3. Create a security NAT destination pool and rule set for the tenant system.

```
[edit tenants tn1 security nat destination]
user@host# set tenants tn1 security nat destination pool h1 address 192.168.1.200
user@host# set tenants tn1 security nat destination rule-set from_zone from zone trust
user@host# set tenants tn1 security nat destination rule-set from_zone rule r1 match source-
address 192.0.2.0/24
user@host# set tenants tn1 security nat destination rule-set from_zone rule r1 match
destination-address 203.0.113.202/24
```

```
user@host# set tenants tn1 security nat destination rule-set from_zone rule r1 then
destination-nat pool h1
```

4. Configure proxy Address Resolution Protocol (ARP).

```
[edit tenant tn1 security nat]
user@host# set tenants tn1 security nat proxy-arp interface ge-0/0/1.0 address 192.0.2.1 to
192.0.2.24
```

Results

From configuration mode, confirm your configuration by entering the `show tenants tn1 security nat` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
source {
    pool pat {
        address {
            192.0.2.1 to 192.0.2.24;
        }
    }
    rule-set from_intf {
        from interface ge-0/0/0.0;
        to interface ge-0/0/1.0;
        rule r1 {
            match {
                source-address 192.168.1.0/24;
                destination-address [203.0.113.200/24 ];
            }
            then {
                source-nat {
                    pool {
                        pat;
                    }
                }
            }
        }
    }
}
destination {
```

```
    pool h1 {
        address 192.168.1.200;
    }
    rule-set from_zone {
        from zone untrust;
        rule r1 {
            match {
                source-address 192.0.2.0/24;
                destination-address 203.0.113.202/24;
            }
            then {
                destination-nat {
                    pool {
                        h1;
                    }
                }
            }
        }
    }
}
static {
    rule-set from_zone {
        from zone untrust;
        rule r1 {
            match {
                source-address 192.0.2.0/24;
                destination-address 203.0.113.203/24;
            }
            then {
                static-nat {
                    prefix {
                        192.168.1.203/24;
                    }
                }
            }
        }
    }
}
proxy-arp {
    interface ge-0/0/1.0 {
        address {
            192.0.2.1 to 192.0.2.24;
        }
```

```
    }
}
```

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

Verification

IN THIS SECTION

- Verifying Static NAT Configuration | 679
- Verifying Destination NAT Configuration | 680
- Verifying Source NAT Configuration | 681

To confirm that the configuration is working properly, perform these tasks:

Verifying Static NAT Configuration

Purpose

To verify that there is traffic matching the static NAT rule set.

Action

From operational mode, enter the `show security nat static rule all tenant tn1` command. View the `Translation hits` field to check for traffic that matches the rule.

```
user@host> show security nat static rule all tenant tn1
```

Sample Output

command-name

```
Total static-nat rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 2/0
Static NAT rule: r1                    Rule-set: from_zone
  Rule-Id                   : 1
```

```
  Rule position             : 1
  From zone                 : untrust
  Source addresses          : 192.0.2.0       - 192.0.2.255
  Destination addresses     : 203.0.113.203
  Host addresses            : 192.168.1.203
  Netmask                   : 32
  Host routing-instance     : N/A
  Translation hits          : 0
    Successful sessions     : 0
    Failed sessions         : 0
  Number of sessions        : 0
```

Meaning

The command output displays the static NAT rule. View the `Translation hits` field to check for traffic that matches the static rule.

Verifying Destination NAT Configuration

Purpose

To verify that there is traffic matching the destination NAT rule set.

Action

From operational mode, enter the `show security nat destination rule all tenant tn1` command. View the `Translation hits` field to check for traffic that matches the rule.

```
user@host> show security nat destination rule all tenant tn1
```

Sample Output

command-name

```
Total destination-nat rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 2/0
Destination NAT rule: r1                    Rule-set: from_zone
  Rule-Id                   : 1
  Rule position             : 1
```

```
  From zone                  : untrust
  Match
    Source addresses         : 192.0.2.0       - 192.0.2.255
    Destination addresses    : 203.0.113.202   - 203.0.113.202
  Action                     : h1
  Translation hits           : 0
    Successful sessions      : 0
    Failed sessions          : 0
  Number of sessions         : 0
```

Meaning

The command output displays the destination NAT rule. View the `Translation hits` field to check for traffic that matches the destination rule.

Verifying Source NAT Configuration

Purpose

To verify that there is traffic matching the source NAT rule set.

Action

From operational mode, enter the `show security nat source rule all tenant tn1` command. View the `Translation hits` field to check for traffic that matches the rule.

```
user@host> show security nat source rule all tenant tn1
```

Sample Output

command-name

```
Total rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 2/0
source NAT rule: r1                    Rule-set: from_intf
  Rule-Id                  : 1
  Rule position            : 1
  From interface           : ge-0/0/0.0
  To interface             : ge-0/0/1.0
```

```
Match
  Source addresses         : 192.168.1.0     - 192.168.1.255
  Destination addresses    : 203.0.113.200   - 203.0.113.200
Action                       : pat
  Persistent NAT type        : N/A
  Persistent NAT mapping type : address-port-mapping
  Inactivity timeout         : 0
  Max session number         : 0
Translation hits           : 0
  Successful sessions      : 0
  Failed sessions          : 0
Number of sessions         : 0
```

Meaning

The command output displays the source NAT rule. View the `Translation hits` field to check for traffic that matches the source rule.

RELATED DOCUMENTATION

Tenant System Configuration Overview | 537

UTM for Tenant Systems

IN THIS SECTION

- Understanding UTM Features in Tenant Systems | 683
- Example: Configuring UTM for the Tenant System | 684

Unified threat management (UTM) provides multiple security features and services for SRX Series devices on the network, protecting users from security threats in a simplified way. UTM secures the tenant systems from viruses, malware, or malicious attachments by scanning the incoming data using Deep Packet Inspection and prevents access to unwanted websites by installing Enhanced Web Filtering (EWF).

Understanding UTM Features in Tenant Systems

Unified Threat Management (UTM) in tenant systems provides several security features such as antispam, antivirus, content filtering, and Web filtering to secure users from multiple Internet-borne threats. The advantage of UTM is streamlined installation and management of these multiple security capabilities. The tenant systems administrator configures the UTM features. Configuring UTM features for tenant systems is similar to configuring UTM features on a device that is not configured for tenant systems.

The security features provided as part of the UTM solution are:

- *Antispam Filtering*—E-mail spam consists of unwanted e-mail messages, usually sent by commercial, malicious, or fraudulent entities. The antispam feature examines transmitted e-mail messages to identify e-mail spam. The default antispam feature is configured at the tenant system administrator and it is applicable for all the tenant systems.

- *Content Filtering*—Content filtering blocks or permits certain types of traffic based on the MIME type, file extension, protocol command, and embedded object type. The default content filtering feature is configured at the tenant system administrator and it is applicable for all the tenant systems.

- *Web Filtering*—Web filtering lets you manage Internet usage by preventing access to inappropriate Web content. The default Web filtering feature is configured at the tenant system administrator, and the tenant system inherit these default Web filtering configuration.

- *Sophos Antivirus* —Sophos Antivirus scanning is offered as a less CPU-intensive alternative to the full file-based antivirus feature. Sophos Antivirus is as an in-the-cloud antivirus solution. The default antivirus feature is configured at the tenant system administrator, and the tenant system inherit these default antivirus configuration.

- *Avira Antivirus* —Avira Antivirus feature profile settings include the scanning options, such as virus detection type, allowlist, blocklist, fallback and notification options. Only one Avira antivirus, Web filtering, Antispam filtering, or Content filtering engine is running in root system. You must configure the Avira antivirus, Web filtering, and Antispam filtering feature type in default configuration. It is configured by the root-user only. All tenants should use the same routing engine and profile type.

You must configure the custom objects for the Web filtering, anti-spam, and content filtering features before configuring the UTM features. You can configure custom objects for each tenant system.

The predefined UTM default policy parameters for Web filtering, content filtering, antivirus, and antispam profiles are configured at the tenant system administrator. The tenant system inherit the same antivirus and Web filtering features configured for the tenant system administrator. The options such as `mime-whitelist` and `url-whitelist` in antivirus profile, and `address-blacklist` and `address-whitelist` in antispam profile can be configured at the following hierarchy levels, respectively:

- `[edit security utm feature-profile anti-virus sophos-engine profile]`

- [edit security utm feature-profile anti-spam sbl profile]

The options url-whitelist and url-blacklist are not supported in the Web fiterling profile, you can use the custom category option to achieve the function.

Example: Configuring UTM for the Tenant System

IN THIS SECTION

- Requirements | 684
- Overview | 684
- Configuration | 685
- Verification | 689

This example shows how to configure the UTM features antivirus, antispam, content filtering, custom message, custom url category, and Web filtering in the tenant system. The tenant system administrator is responsible for assigning the UTM features to the tenant system.

Requirements

This example uses the following hardware and software components:

- SRX Series device configured with the tenant systems.
- Junos OS Release 19.2R1 and later releases.

Before you begin:

- Understand the tenant systems role and functions. See tenant systems overview.

Overview

The tenant system administrator assigns UTM features antivirus, antispam, content filtering, and Web filtering to the tenant system.

This example shows how to configure the UTM features for tenant system.

Configuration

IN THIS SECTION

- CLI Quick Configuration | 685
- Configuring UTM for Tenant System | 686
- Results | 687

CLI Quick Configuration

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

```
set tenants TSYS1 security utm custom-objects url-pattern cust-list value www.ask.com
set tenants TSYS1 security utm custom-objects url-pattern cust-list value www.playboy.com
set tenants TSYS1 security utm custom-objects url-pattern cust-list2 value www.baidu.com
set tenants TSYS1 security utm custom-objects custom-url-category cust-list value cust-list
set tenants TSYS1 security utm custom-objects custom-url-category cust-list2 value cust-list2
set tenants TSYS1 security utm feature-profile web-filtering juniper-local profile my_local1
default log-and-permit
set tenants TSYS1 security utm feature-profile web-filtering juniper-local profile my_local1
category cust-list action log-and-permit
set tenants TSYS1 security utm feature-profile web-filtering juniper-local profile my_local1
category cust-list2 action block
set tenants TSYS1 security utm feature-profile web-filtering juniper-local profile my_local1
fallback-settings default log-and-permit
set tenants TSYS1 security utm feature-profile web-filtering juniper-enhanced profile
ewf_my_profile1 category Enhanced_Adult_Content action block
set tenants TSYS1 security utm feature-profile web-filtering juniper-enhanced profile
ewf_my_profile1 category Enhanced_Social_Web_Facebook action log-and-permit
set tenants TSYS1 security utm feature-profile web-filtering juniper-enhanced profile
ewf_my_profile1 category cust-list action block
set tenants TSYS1 security utm utm-policy utmpolicy1 web-filtering http-profile ewf_my_profile1
```

Configuring UTM for Tenant System

Step-by-Step Procedure

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

1. Log in to the tenant system and enter configuration mode.

```
user@host> configure
admin@host#
```

2. Configure the custom objects for the tenant system.

```
[edit tenants TSYS1 security utm custom-objects]
user@host# url-pattern cust-list value www.ask.com
user@host# url-pattern cust-list value www.playboy.com
user@host# url-pattern cust-list2 value www.baidu.com
user@host# custom-url-category cust-list value cust-list
user@host# custom-url-category cust-list2 value cust-list2
```

3. Configure the feature profile `web-filtering` for the tenant system.

```
[edit tenants TSYS1 security utm feature-profile]
user@host# set tenants TSYS1 security utm feature-profile web-filtering juniper-local profile
my_local1 default log-and-permit
user@host# set tenants TSYS1 security utm feature-profile web-filtering juniper-local profile
my_local1 category cust-list action log-and-permit
user@host# set tenants TSYS1 security utm feature-profile web-filtering juniper-local profile
my_local1 category cust-list2 action block
user@host# set tenants TSYS1 security utm feature-profile web-filtering juniper-local profile
my_local1 fallback-settings default log-and-permit
user@host# set tenants TSYS1 security utm feature-profile web-filtering juniper-enhanced
profile ewf_my_profile1 category Enhanced_Adult_Content action block
user@host# set tenants TSYS1 security utm feature-profile web-filtering juniper-enhanced
profile ewf_my_profile1 category Enhanced_Social_Web_Facebook action log-and-permit
user@host# set tenants TSYS1 security utm feature-profile web-filtering juniper-enhanced
profile ewf_my_profile1 category cust-list action block
```

4. Configure the UTM policy for the tenant system.

```
[edit tenants TSYS1 security utm ]
user@host# set tenants TSYS1 security utm utm-policy utmpolicy1 web-filtering http-profile
ewf_my_profile1
```

Results

- From configuration mode, confirm your configuration by entering the `show tenants TSYS1 security utm custom-objects` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show tenants TSYS1 security utm custom-objects
url-pattern {
    cust-list {
        value [ www.ask.com www.playboy.com ];
    }
    cust-list2 {
        value www.baidu.com;
    }
}
custom-url-category {
    cust-list {
        value cust-list;
    }
    cust-list2 {
        value cust-list2;
    }
}
```

- From configuration mode, confirm your configuration by entering the `show tenants TSYS1 security utm feature-profile web-filtering` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show tenants TSYS1 security utm feature-profile web-filtering
juniper-local {
    profile my_local1 {
        default log-and-permit;
        category {
            cust-list {
```

```
                action log-and-permit;
            }
            cust-list2 {
                action block;
            }
        }
        fallback-settings {
            default log-and-permit;
        }
    }
}
juniper-enhanced {
    profile ewf_my_profile1 {
        category {
            Enhanced_Adult_Content {
                action block;
            }
            Enhanced_Social_Web_Facebook {
                action log-and-permit;
            }
            cust-list {
                action block;
            }
        }
    }
}
```

- From configuration mode, confirm your configuration by entering the `show tenants TSYS1 security utm` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show tenants TSYS1 security utm
utm-policy utmpolicy1 {
    web-filtering {
        http-profile ewf_my_profile1;
    }
}
```

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

Verification

IN THIS SECTION

- Verifying Web Filtering Configuration | 689

To confirm that the configuration is working properly, perform these tasks:

Verifying Web Filtering Configuration

Purpose

Verify that the Web filtering feature is configured for the tenant system.

Action

From operational mode, enter the `show security utm web-filtering statistics tenant TSYS1` command to view the details of the Web filtering feature configured for the tenant system.

```
user@host> show security utm web-filtering statistics tenant TSYS1
UTM web-filtering statistics:
    Total requests:                   19784932
    white list hit:                   0
    Black list hit:                   0
    No license permit:                0
    Queries to server:                19782736
    Server reply permit:              18819472
    Server reply block:               0
```

Meaning

The output displays the Web filtering statistics for the tenant system.

IDP for Tenant Systems

IN THIS SECTION

- Understanding IDP for Tenant Systems | 690
- Understanding IDP Features in Tenant Systems | 692
- Example: Configuring IDP Policies and Attacks for Tenant Systems | 694

An Intrusion Detection and Prevention (IDP) policy in tenant systems enables you to selectively enforce various attack detection and prevention techniques on the network traffic passing through an SRX Series device. The SRX Series devices offer the same set of IDP signatures that are available on Juniper Networks IDP Series Intrusion Detection and Prevention Appliances to secure networks against attacks.

Understanding IDP for Tenant Systems

IN THIS SECTION

- IDP Policies | 690
- Limitation | 692
- IDP Installation and Licensing for Tenant Systems | 692

A Junos OS Intrusion Detection and Prevention (IDP) policy enables you to selectively enforce various attack detection and prevention techniques on network traffic passing through a tenant system.

This topic includes the following sections:

IDP Policies

Configuring IDP policies at the root level and tenant systems level are similar. IDP policy templates configured at the root level are visible and used by all tenant systems. The primary administrator specifies an IDP policy in the security profile that is bound to a tenant system. To enable IDP in a tenant

system, the primary administrator or tenant system administrator configures a security policy that defines the traffic to be inspected and specifies at the `permit application-services idp-policy idp-policy-name` hierarchy level.

The primary administrator can configure multiple IDP policies and a tenant system can have multiple IDP policies at a time. For tenant systems, the primary administrator can either bind the same IDP policy to multiple tenant systems or bind the necessary IDP policies to each tenant system. If you configure more than one IDP policy, then configuring a default IDP policy is mandatory.

The primary administrator configures the number of maximum IDP sessions reservation for a primary logical system and tenant systems. The number of IDP sessions that are allowed for a primary logical system are defined using the command `set security idp max-sessions max-sessions` and the number of IDP sessions that are allowed for a tenant system are defined using the command `set security idp tenant-system tenant-system max-sessions max-sessions`.

The tenant system administrator performs the following actions:

- Configure multiple IDP policies and attach to the firewall policies to be used by the tenant systems. If the IDP policy is not configured for a tenant system, the default IDP policy configured by the primary administrator is used. The IDP policy is bound to the tenant systems through a tenant systems security policy.
- Create or modify IDP policies for their tenant system. The IDP policies are bound to tenant systems. When an IDP policy is changed, and commit fails, only the tenant system that has initiated the commit change is notified about the commit failure.
- The tenant system administrator can create security zones in the tenant system and assign interfaces to each security zone. Zones that are specific to tenant systems cannot be referenced in IDP policies configured by the primary administrator. The primary administrator can reference zones in the primary logical system in an IDP policy configured for the primary logical system.
- View the attack statistics detected and IDP counters, attack table, and policy commit status by the individual tenant system using the commands `show security idp counters`, `show security idp attack table`, `show security idp policies`, `show security idp policy-commit-status`, and `show security idp security-package-version`.

View the attack statistics detected and IDP counters, attack table, and policy commit status from the root using the commands `show security idp counters counters tenant tenant-name`, `show security idp attack table tenant tenant-name`, `show security idp policies tenant tenant-name`, `show security idp policy-commit-status tenant tenant-name`, and `show security idp security-package-version tenant tenant-name`.

Limitation

- IDP policy compilation in Packet Forwarding Engine is done at global level. Any changes in policy made for a logical system or a tenant system results in the compilation of policies of all the logical systems or tenant systems because the IDP internally treats it as a single global policy.
- Any changes in policy made for a logical system or a tenant system results in clearing the attack table of all logical systems or a tenant systems.

IDP Installation and Licensing for Tenant Systems

An idp-sig license must be installed at the root level. Once IDP is enabled at the root level, it can be used with any tenant system on the device.

A single IDP security package is installed for all tenant systems on the device at the root level. The download and install options can only be executed at the root level. The same version of the IDP attack database is shared by all tenant systems.

Understanding IDP Features in Tenant Systems

IN THIS SECTION

- Rulebases | **692**
- Multi-Detectors | **693**
- Logging and Monitoring | **693**

This topic includes the following sections:

Rulebases

A single IDP policy can contain only one instance of any type of rulebase. The Intrusion prevention system (IPS) rulebase uses attack objects to detect known and unknown attacks. It detects attacks based on stateful signature and protocol anomalies.

NOTE: Status monitoring for IPS is global to the device and not on a per tenant system basis.

Multi-Detectors

When a new IDP security package is received, it contains attack definitions and a detector. After a new policy is loaded, it is also associated with a detector. If the policy being loaded has an associated detector that matches the detector already in use by the existing policy, the new detector is not loaded and both policies use a single associated detector. But if the new detector does not match the current detector, the new detector is loaded along with the new policy. In this case, each loaded policy will then use its own associated detector for attack detection.

The version of the detector is common to all tenant systems.

Logging and Monitoring

Status monitoring options are available to the primary administrator only. All status monitoring options under the `show security idp` and `clear security idp` CLI operational commands present global information, but not on a per tenant system basis.

NOTE:

- SNMP monitoring for IDP is not supported on tenant systems.
- The tenant systems supports only the stream mode for syslog and does not support the event mode.

IDP generates event logs when an event matches an IDP policy rule in which logging is enabled.

The tenant systems identification is added to the following types of IDP traffic processing logs:

- Attack logs. The following example shows an attack log for the `TSYS1` tenant system:

```
"<14>1 2019-02-18T02:17:56+05:30 4.0.0.254 pamba RT_IDP  - -  IDP_ATTACK_LOG_EVENT_LS: Lsys
TSYS1: IDP: At 1550485076, SIG Attack log <4.0.0.1/51480->5.0.0.1/21> for TCP protocol and
service SERVICE_IDP application FTP by rule 1 of rulebase IPS in policy new. attack: id=4641,
repeat=0, action=NONE, threat-severity=MEDIUM, name=FTP:USER:ROOT, NAT <0.0.0.0:0-
>0.0.0.0:0>, time-elapsed=0, inbytes=0, outbytes=0, inpackets=0, outpackets=0,
```

```
intf:l1z1:xe-4/0/0.0->l1z2:xe-4/0/1.0, packet-log-id: 0, alert=no, username=N/A, roles=N/A
and misc-message -
```

- IP action logs. The following example shows an IP action log for the TSYS1 tenant system:

```
"<14>1 2019-02-19T02:21:43+05:30 4.0.0.254 pamba RT_FLOW  - -  FLOW_IP_ACTION_LS: Lsys TSYS1:
Flow IP action detected attack attempt:4.0.0.1/51492 --> 5.0.0.1/21 from interface xe
-4/0/0.0, from zone l1z1, action close.
"<14>1 2019-02-19T02:21:45+05:30 4.0.0.254 pamba RT_FLOW  - -  APPTRACK_SESSION_CLOSE_LS:
Lsys TSYS1: AppTrack session closed Closed by junos-tcp-clt-emul: 4.0.0.1/51492-
>5.0.0.1/        21 junos-ftp FTP UNKNOWN 4.0.0.1/51492->5.0.0.1/21 N/A N/A 6 l1z1-l1z2 l1z1
l1z2 50000058 6(287) 5(281) 6 N/A N/A No N/A N/A VR1 xe-4/0/1.0  0 0 Infrastructure File-
Servers N/A         N/A
```

Example: Configuring IDP Policies and Attacks for Tenant Systems

IN THIS SECTION

- Requirements | 694
- Overview | 695
- Configuration | 695
- Verification | 709

This example shows how to configure IDP policies and attacks for tenant systems.

Requirements

This example uses the following hardware and software components:

- SRX Series device configured with the tenant systems.
- Junos OS Release 19.2R1 and later releases.

Before you configure IDP policies and attacks for tenant systems, be sure you have:

- Read "Tenant Systems Overview" on page 530 to understand how this task fits into the overall configuration process.

- Create tenant system TSYS1. See Example: Creating Tenant Systems, Tenant System Administrators, and an Interconnect VPLS Switch.
- Create security zones for tenant system TSYS1. See "Example: Configuring Zones in the Tenant System" on page 574.
- Log in to the tenant system as the tenant system administrator. See "Tenant System Configuration Overview" on page 537.

Overview

In this example you configure IDP custom attacks, policies, custom attack group, pre-defined attack and attack-group, and dynamic attack group in the tenant system TSYS1.

Configuration

IN THIS SECTION

- Configuring a Custom Attack | **695**
- Configuring an IDP Policy | **697**
- Configuring Multiple IDP Policies with a Default IDP Policy | **699**
- Configuring IDP Custom Attack Group | **703**
- Configuring Pre-defined Attack and Attack Group | **706**
- Configuring IDP Dynamic Attack Group | **708**

Configuring a Custom Attack

CLI Quick Configuration

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

```
set security idp custom-attack my-http severity info
set security idp custom-attack my-http attack-type signature protocol-binding application HTTP
set security idp custom-attack my-http attack-type signature context http-get-url
set security idp custom-attack my-http attack-type signature pattern .*testing.*
set security idp custom-attack my-http attack-type signature direction any
```

Step-by-Step Procedure

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

To configure a custom attack object:

1. Create the custom attack object and set the severity level.

```
[edit security idp]
user@host:TSYS1# set custom-attack my-http severity info
```

2. Configure stateful signature parameters.

```
[edit security idp]
user@host:TSYS1# set custom-attack my-http attack-type signature protocol-binding application
HTTP
user@host:TSYS1# set custom-attack my-http attack-type signature context http-get-url
user@host:TSYS1# set custom-attack my-http attack-type signature pattern .*testing.*
user@host:TSYS1# set custom-attack my-http attack-type signature direction any
```

Results

From configuration mode, confirm your configuration by entering the `show security idp custom-attack my-http` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host:TSYS1# show security idp custom-attack my-http
severity info;
    attack-type {
    signature {
        protocol-binding {
            application HTTP;
        }
        context http-get-url;
        pattern .*testing.*;
        direction any;
```

```
    }
}
```

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

Configuring an IDP Policy

CLI Quick Configuration

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

```
set security idp idp-policy idpengine rulebase-ips rule 1 match from-zone any
set security idp idp-policy idpengine rulebase-ips rule 1 match source-address any
set security idp idp-policy idpengine rulebase-ips rule 1 match to-zone any
set security idp idp-policy idpengine rulebase-ips rule 1 match destination-address any
set security idp idp-policy idpengine rulebase-ips rule 1 match application default
set security idp idp-policy idpengine rulebase-ips rule 1 match attacks custom-attacks my-http
set security idp idp-policy idpengine rulebase-ips rule 1 then action no-action
set security idp idp-policy idpengine rulebase-ips rule 1 then notification log-attacks
```

Step-by-Step Procedure

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

To configure an IDP policy:

1. Create the IDP policy and configure match conditions.

```
[edit security idp]
user@host:TSYS1# set idp-policy idpengine rulebase-ips rule 1 match from-zone any
user@host:TSYS1# set idp-policy idpengine rulebase-ips rule 1 match source-address any
user@host:TSYS1# set idp-policy idpengine rulebase-ips rule 1 match to-zone any
user@host:TSYS1# set idp-policy idpengine rulebase-ips rule 1 match destination-address any
user@host:TSYS1# set idp-policy idpengine rulebase-ips rule 1 match application default
user@host:TSYS1# set idp-policy idpengine rulebase-ips rule 1 match attacks custom-attacks my-
http
```

2. Configure actions for the IDP policy.

```
[edit security idp]
user@host:TSYS1# set idp-policy idpengine rulebase-ips rule 1 then action no-action
user@host:TSYS1# set idp-policy idpengine rulebase-ips rule 1 then notification log-attacks
```

Results

From configuration mode, confirm your configuration by entering the `show security idp idp-policy idpengine` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host:TSYS1# show security idp idp-policy idpengine
rulebase-ips {
    rule 1 {
        match {
            from-zone any;
            source-address any;
            to-zone any;
            destination-address any;
            application default;
            attacks {
                custom-attacks my-http;
            }
        }
        then {
            action {
                no-action;
            }
            notification {
                log-attacks;
            }
        }
    }
}
```

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

Configuring Multiple IDP Policies with a Default IDP Policy

CLI Quick Configuration

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

```
set security idp idp-policy idpengine rulebase-ips rule 1 match from-zone any
set security idp idp-policy idpengine rulebase-ips rule 1 match source-address any
set security idp idp-policy idpengine rulebase-ips rule 1 match to-zone any
set security idp idp-policy idpengine rulebase-ips rule 1 match destination-address any
set security idp idp-policy idpengine rulebase-ips rule 1 match application default
set security idp idp-policy idpengine rulebase-ips rule 1 match attacks predefined-attacks
HTTP:AUDIT:URL
set security idp idp-policy idpengine rulebase-ips rule 1 then action no-action
set security idp idp-policy idpengine rulebase-ips rule 1 then notification log-attacks
set security idp idp-policy idpengine1 rulebase-ips rule 1 match from-zone any
set security idp idp-policy idpengine1 rulebase-ips rule 1 match source-address any
set security idp idp-policy idpengine1 rulebase-ips rule 1 match to-zone any
set security idp idp-policy idpengine1 rulebase-ips rule 1 match destination-address any
set security idp idp-policy idpengine1 rulebase-ips rule 1 match attacks predefined-attacks
FTP:USER:ROOT
set security idp idp-policy idpengine1 rulebase-ips rule 1 then action no-action
set security idp idp-policy idpengine1 rulebase-ips rule 1 then notification log-attacks
set security policies from-zone l1z1 to-zone l1z2 policy l1z1-l1z2 match source-address any
set security policies from-zone l1z1 to-zone l1z2 policy l1z1-l1z2 match destination-address any
set security policies from-zone l1z1 to-zone l1z2 policy l1z1-l1z2 match application any
set security policies from-zone l1z1 to-zone l1z2 policy l1z1-l1z2 match dynamic-application
junos:FTP
set security policies from-zone l1z1 to-zone l1z2 policy l1z1-l1z2 then permit application-
services idp-policy idpengine1
set security policies from-zone l1z1 to-zone l1z2 policy 2 match source-address any
set security policies from-zone l1z1 to-zone l1z2 policy 2 match destination-address any
set security policies from-zone l1z1 to-zone l1z2 policy 2 match application any
set security policies from-zone l1z1 to-zone l1z2 policy 2 match dynamic-application junos:HTTP
set security policies from-zone l1z1 to-zone l1z2 policy 2 then permit application-services idp-
policy idpengine
set security idp default-policy idpengine1
```

Step-by-Step Procedure

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

To configure multiple IDP policies:

1. Create multiple IDP policies and configure match conditions.

```
[edit security idp]
user@host:TSYS1# set idp-policy idpengine rulebase-ips rule 1 match from-zone any
user@host:TSYS1# set idp-policy idpengine rulebase-ips rule 1 match source-address any
user@host:TSYS1# set idp-policy idpengine rulebase-ips rule 1 match to-zone any
user@host:TSYS1# set idp-policy idpengine rulebase-ips rule 1 match destination-address any
user@host:TSYS1# set idp-policy idpengine rulebase-ips rule 1 match application default
user@host:TSYS1# set idp-policy idpengine rulebase-ips rule 1 match attacks predefined-
attacks HTTP:AUDIT:URL
user@host:TSYS1# set idp-policy idpengine rulebase-ips rule 1 then action no-action
user@host:TSYS1# set idp-policy idpengine rulebase-ips rule 1 then notification log-attacks
user@host:TSYS1# set idp-policy idpengine1 rulebase-ips rule 1 match from-zone any
user@host:TSYS1# set idp-policy idpengine1 rulebase-ips rule 1 match source-address any
user@host:TSYS1# set idp-policy idpengine1 rulebase-ips rule 1 match to-zone any
user@host:TSYS1# set idp-policy idpengine1 rulebase-ips rule 1 match destination-address any
user@host:TSYS1# set idp-policy idpengine1 rulebase-ips rule 1 match attacks predefined-
attacks FTP:USER:ROOT
user@host:TSYS1# set idp-policy idpengine1 rulebase-ips rule 1 then action no-action
user@host:TSYS1# set idp-policy idpengine1 rulebase-ips rule 1 then notification log-attacks
```

2. Configure security policies and attach IDP policies to them.

```
[edit security policies]
user@host:TSYS1# set from-zone l1z1 to-zone l1z2 policy l1z1-l1z2 match source-address any
user@host:TSYS1# set from-zone l1z1 to-zone l1z2 policy l1z1-l1z2 match destination-address
any
user@host:TSYS1# set from-zone l1z1 to-zone l1z2 policy l1z1-l1z2 match application any
user@host:TSYS1# set from-zone l1z1 to-zone l1z2 policy l1z1-l1z2 match dynamic-application
junos:FTP
user@host:TSYS1# set from-zone l1z1 to-zone l1z2 policy l1z1-l1z2 then permit application-
services idp-policy idpengine1
user@host:TSYS1# set from-zone l1z1 to-zone l1z2 policy 2 match source-address any
user@host:TSYS1# set from-zone l1z1 to-zone l1z2 policy 2 match destination-address any
```

```
user@host:TSYS1# set from-zone l1z1 to-zone l1z2 policy 2 match application any
user@host:TSYS1# set from-zone l1z1 to-zone l1z2 policy 2 match dynamic-application junos:HTTP
user@host:TSYS1# set from-zone l1z1 to-zone l1z2 policy 2 then permit application-services
idp-policy idpengine
```

3. Configure a default IDP policy.

> **NOTE**: If you configure more than one IDP policy, then configuring a default IDP policy is mandatory.

```
[edit security idp]
user@host:TSYS1# set default-policy idpengine1
```

Results

From configuration mode, confirm your configuration by entering the `show security idp idp-policy idpengine`, `show security idp idp-policy idpengine1`, `show security policies`, and `show security policies` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host:TSYS1# show security idp idp-policy idpengine
rulebase-ips {
    rule 1 {
        match {
            from-zone any;
            source-address any;
            to-zone any;
            destination-address any;
            application default;
            attacks {
                predefined-attacks HTTP:AUDIT:URL;
            }
        }
        then {
            action {
                no-action;
            }
```

```
            notification {
                log-attacks;
            }
        }
    }
}
```

```
[edit]
user@host:TSYS1# show security idp idp-policy idpengine1
rulebase-ips {
    rule 1 {
        match {
            from-zone any;
            source-address any;
            to-zone any;
            destination-address any;
            attacks {
                predefined-attacks FTP:USER:ROOT;
            }
        }
        then {
            action {
                no-action;
            }
            notification {
                log-attacks;
            }
        }
    }
}
```

```
[edit]
user@host:TSYS1# show security policies
from-zone l1z1 to-zone l1z2 {
    policy l1z1-l1z2 {
        match {
            source-address any;
            destination-address any;
            application any;
            dynamic-application junos:FTP;
```

```
        }
        then {
            permit {
                application-services {
                    idp-policy idpengine1;
                }
            }
        }
    }
    policy 2 {
        match {
            source-address any;
            destination-address any;
            application any;
            dynamic-application junos:HTTP;
        }
        then {
            permit {
                application-services {
                    idp-policy idpengine;
                }
            }
        }
    }
}
```

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

Configuring IDP Custom Attack Group

CLI Quick Configuration

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

```
set security idp idp-policy idpengine rulebase-ips rule 1 match attacks custom-attack-groups
cust-group
set security idp idp-policy idpengine rulebase-ips rule 1 then action no-action
set security idp idp-policy idpengine rulebase-ips rule 1 then notification log-attacks
set security idp custom-attack customftp severity warning
set security idp custom-attack customftp attack-type signature context ftp-username
```

```
set security idp custom-attack customftp attack-type signature pattern .*guest.*
set security idp custom-attack customftp attack-type signature direction client-to-server
set security idp custom-attack-group cust-group group-members customftp
set security idp custom-attack-group cust-group group-members ICMP:INFO:TIMESTAMP
set security idp custom-attack-group cust-group group-members "FTP - Minor"
set security idp custom-attack-group cust-group group-members dyn1
set security idp dynamic-attack-group dyn1 filters category values HTTP
```

Step-by-Step Procedure

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

To configure IDP custom attack group:

1. Create the IDP policy.

```
[edit security idp]
user@host:TSYS1# set idp-policy idpengine rulebase-ips rule 1 match attacks custom-attack-
groups cust-group
```

2. Configure match condition of IDP policy.

```
[edit security idp]
user@host:TSYS1# set security idp idp-policy idpengine rulebase-ips rule 1 then action no-
action
user@host:TSYS1# set security idp idp-policy idpengine rulebase-ips rule 1 then notification
log-attacks
```

3. Configure stateful signature parameters.

```
[edit security idp]
user@host:TSYS1# set security idp custom-attack customftp severity warning
user@host:TSYS1# set custom-attack customftp attack-type signature context ftp-username
user@host:TSYS1# set custom-attack customftp attack-type signature pattern .*guest.*
user@host:TSYS1# set custom-attack customftp attack-type signature direction client-to-server
user@host:TSYS1# set custom-attack-group cust-group group-members customftp
user@host:TSYS1# set custom-attack-group cust-group group-members ICMP:INFO:TIMESTAMP
user@host:TSYS1# set custom-attack-group cust-group group-members "FTP - Minor"
```

```
user@host:TSYS1# set custom-attack-group cust-group group-members dyn1
user@host:TSYS1# set dynamic-attack-group dyn1 filters category values HTTP
```

Results

From configuration mode, confirm your configuration by entering the `show security idp` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host:TSYS1# show security idp
idp-policy idpengine {
    rulebase-ips {
        rule 1 {
            match {
                attacks {
                    custom-attack-groups cust-group;
                }
            }
            then {
                action {
                    no-action;
                }
                notification {
                    log-attacks;
                }
            }
        }
    }
}
custom-attack customftp {
    severity warning;
    attack-type {
        signature {
            context ftp-username;
            pattern .*guest.*;
            direction client-to-server;
        }
    }
}
custom-attack-group cust-group {
```

```
    group-members [ customftp ICMP:INFO:TIMESTAMP "FTP - Minor" dyn1 ];
}
dynamic-attack-group dyn1 {
    filters {
        category {
            values HTTP;
        }
    }
}
```

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

Configuring Pre-defined Attack and Attack Group

CLI Quick Configuration

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

```
set security idp idp-policy idpengine rulebase-ips rule 1 match attacks predefined-attacks
FTP:USER:ROOT
set security idp idp-policy idpengine rulebase-ips rule 1 match attacks predefined-attack-groups
"HTTP - All"
```

Step-by-Step Procedure

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

To configure the pre-defined attack and attack group:

1. Configure the pre-defined attack.

```
[edit security idp]
user@host:TSYS1# set idp-policy idpengine rulebase-ips rule 1 match attacks predefined-
attacks FTP:USER:ROOT
```

2. Configure the pre-defined attack group.

```
[edit security idp]
user@host:TSYS1# set idp-policy idpengine rulebase-ips rule 1 match attacks predefined-attack-
groups "HTTP - All"
```

Results

From configuration mode, confirm your configuration by entering the `show security idp idp-policy idpengine` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host:TSYS1# show security idp idp-policy idpengine
rulebase-ips {
    rule 1 {
        match {
            attacks {
                predefined-attacks FTP:USER:ROOT;
                predefined-attack-groups "HTTP - All";
            }
        }
        then {
            action {
                no-action;
            }
            notification {
                log-attacks;
            }
        }
    }
}
```

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

Configuring IDP Dynamic Attack Group

CLI Quick Configuration

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

```
set security idp dynamic-attack-group dyn1 filters direction values server-to-client
```

Step-by-Step Procedure

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

To configure IDP dynamic attack group:

1. Configure dynamic attack group parameter.

```
[edit security idp]
user@host:TSYS1# set dynamic-attack-group dyn1 filters direction values server-to-client
```

Results

From configuration mode, confirm your configuration by entering the `show security idp` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host:TSYS1# show security idp
dynamic-attack-group dyn1 {
    filters {
        direction {
            values server-to-client;
        }
    }
}
```

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

Verification

IN THIS SECTION

- Verify IDP Policies and Commit Status | 709
- Verify IDP Attack Detection | 710
- Verify IDP Counters | 710

Verify IDP Policies and Commit Status

Purpose

Verify that the IDP policies and commit status is displayed after policy compilation for the tenant system TSYS1.

Action

From operational mode, enter the `show security idp policies` command.

```
user@host:TSYS1> show security idp policies
ID    Name                  Sessions    Memory      Detector
 1     idpengine             0           186024      12.6.130180122
```

From operational mode, enter the `show security idp policy-commit-status` command.

```
user@host:TSYS1> show security idp policy-commit-statusIDP policy[/var/db/idpd/bins//idp-policy-
unified.bin.gz.v] and detector[/var/db/idpd/sec-repository/installed-detector/libidp-
detector.so.tgz.v] loaded successfully.
The loaded policy size is:2912 Bytes
```

Meaning

The output displays the IDP policy configured in the tenant system TSYS1 and the commit status information.

Verify IDP Attack Detection

Purpose

Verify that the IDP attack detection is successful for the tenant system TSYS1 and displayed in the attack table.

Action

From operational mode, enter the show security idp attack table command.

```
user@host:TSYS1> show security idp attack table
IDP attack statistics:
  Attack name                                #Hits
  my-http                                        1
```

Meaning

The output displays the attacks detected for the custom attack that is configured in the tenant system TSYS1.

Verify IDP Counters

Purpose

Verify one of the IDP counter status is displayed for the tenant system TSYS1.

Action

From operational mode, enter the show security idp counters flow command.

```
user@host:TSYS1> show security idp counters flow
IDP counters:

  IDP counter type                                                  Value
Fast-path packets                                                   38
Slow-path packets                                                   1
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                                                 1
Policy cache entries                                                0
Maximum flow hash collisions                                        0
Flow hash collisions                                                0
Gates added                                                         0
Gate matches                                                        0
Sessions deleted                                                    1
Sessions aged-out                                                   0
Sessions in-use while aged-out                                      0
TCP flows marked dead on RST/FIN                                    1
Policy init failed                                                  0
Policy reinit 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
IDP session gate creation requests                                  0
IDP session gate creation acknowledgements                          0
IDP session gate hits                                               0
IDP session gate timeouts                                           0
Number of times Sessions crossed the CPU threshold value that is set    0
Number of times Sessions crossed the CPU upper threshold            0
Sessions constructed                                                1
SM Sessions ignored                                                 0
SM Sessions dropped                                                 0
SM Sessions interested                                              2
SM Sessions not interested                                          0
SM Sessions interest error                                          0
Sessions destructed                                                 1
SM Session Create                                                   1
SM Packet Process                                                   38
SM ftp data session ignored by idp                                  1
SM Session close                                                    1
SM Client-to-server packets                                         15
SM Server-to-client packets                                         23
SM Client-to-server L7 bytes                                        99
```

SM Server-to-client L7 bytes	367
Client-to-server flows ignored	0
Server-to-client flows ignored	0
Server-to-client flows tcp optimized	0
Client-to-server flows tcp optimized	0
Both directions flows ignored	1
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
NGAppID Events with no L7 App	0
NGAppID Events with no active-policy	0
NGAppID Detector failed from event handler	0
NGAppID Detector failed from API	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	37
kq op error	0

```
kq op stop                                                          0
PME wait not set                                                    0
PME wait set                                                        0
PME KQ run not called                                               0
IDP sessions ignored for content decompression in intel inspect mode    0
IDP sessions ignored for bytes depth limit in intel inspect mode        0
IDP sessions ignored for protocol decoding in intel inspect mode        0
IDP sessions detected CPU usage crossed intel inspect CPU threshold     0
IDP sessions detected mem drop below intel inspect low mem threshold    0
```

Meaning

The output displays the IDP counter flow status is displayed properly for the tenant system TSYS1.

ALG for Tenant Systems

IN THIS SECTION

- Understanding ALG Support for Tenant System | **713**
- Enabling and Disabling ALG for Tenant System | **714**
- Example: Configuring ALG in Tenant System | **719**

An Application Layer Gateway (ALG) in tenant systems enables the gateway to parse application layer payloads and take decisions whether to allow or deny traffic to the application server. ALGs supports the applications such as Transfer Protocol (FTP) and various IP protocols that use the application layer payload to communicate the dynamic Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) ports on which the applications open data connections. For more information, see the following topics:

Understanding ALG Support for Tenant System

An Application Layer Gateway (ALG) enables the gateway to parse application layer payloads and take decisions whether to allow or deny traffic to the application server.

Starting in Junos OS Release 18.3R1, the ALG feature supported on logical systems is now extended on tenants systems.

The tenant systems administrator can configure the ALG features for the tenant systems. The primary administrator can configure the ALG features and display the ALG information for all tenants. The tenant systems administrator can only apply configurations and display information in its own tenant.

Each tenant system displays the ALG counters to monitor the traffic. For example, use commands `show security alg sip counters tenants TN1` to get SIP counters in tenant systems and `show security alg sip counters tenants all` to get SIP counters in all existing tenant systems.

Enabling the security log for the tenant generates the ALG logs per tenant.

NOTE: When you upgrade to Junos OS Release 18.3R1, the ALG status for each tenant system might be different depending on the default configuration or configuration in a release prior to Junos OS Release 18.3R1. We recommend you to change the ALG configurations for tenant systems as per your requirements after an upgrade to latest Junos OS version.

Enabling and Disabling ALG for Tenant System

This topic shows how to enable or disable the ALG status for each tenant system.

1. By Default IKE ALG is disabled on the tenant system. To enable this ALG, use the following command.
   - Enable IKE and ESP ALG with NAT.

```
[edit]
user@host# set tenants TN1 security alg ike-esp-nat enable
```

2. By default, the DNS, FTP, PPTP, SIP, SUNRPC and TWAMP ALGs are enabled on the tenant system. To disable these ALGs, use the following commands.
   - Disable DNS ALG.

```
[edit]
user@host# set tenants TN1 security alg dns disable
```

- Disable FTP ALG.

```
[edit]
user@host# set tenants TN1 security alg ftp disable
```

- Disable H323 ALG.

```
[edit]
user@host# set tenants TN1 security alg h323 disable
```

- Disable MGCP ALG.

```
[edit]
user@host# set tenants TN1 security alg mgcp disable
```

- Disable MSRPC ALG.

```
[edit]
user@host# set tenants TN1 security alg msrpc disable
```

- Disable PPTP ALG.

```
[edit]
user@host# set tenants TN1 security alg pptp disable
```

- Disable RSH ALG.

```
[edit]
user@host# set tenants TN1 security alg rsh disable
```

- Disable RTSP ALG.

```
[edit]
user@host# set tenants TN1 security alg rtsp disable
```

- Disable SCCP ALG.

```
[edit]
user@host# set tenants TN1 security alg sccp disable
```

- Disable SIP ALG.

```
[edit]
user@host# set tenants TN1 security alg sip disable
```

- Disable SQL ALG.

```
[edit]
user@host# set tenants TN1 security alg sql disable
```

- Disable SUNRPC ALG.

```
[edit]
user@host# set tenants TN1 security alg sunrpc disable
```

- Disable TALK ALG.

```
[edit]
user@host# set tenants TN1 security alg talk disable
```

- Disable TFTP ALG.

```
[edit]
user@host# set tenants TN1 security alg tftp disable
```

3. Configuring ALG functions in tenant systems.

- Configure DNS ALG.

```
[edit]
user@host# set tenants TN1 security alg dns
```

- Configure FTP ALG.

```
[edit]
user@host# set tenants TN1 security alg ftp
```

- Configure H323 ALG.

```
[edit]
user@host# set tenants TN1 security alg h323
```

- Configure IKE and ESP ALG with NAT.

```
[edit]
user@host# set tenants TN1 security alg ike-esp-nat
```

- Configure MGCP ALG.

```
[edit]
user@host# set tenants TN1 security alg mgcp
```

- Configure MSRPC ALG.

```
[edit]
user@host# set tenants TN1 security alg msrpc
```

- Configure PPTP ALG.

```
[edit]
user@host# set tenants TN1 security alg pptp
```

- Configure RSH ALG.

```
[edit]
user@host# set tenants TN1 security alg rsh
```

- Configure RTSP ALG.

```
[edit]
user@host# set tenants TN1 security alg rtsp
```

- Configure SCCP ALG.

```
[edit]
user@host# set tenants TN1 security alg sccp
```

- Configure SIP ALG.

```
[edit]
user@host# set tenants TN1 security alg sip
```

- Configure SQL ALG.

```
[edit]
user@host# set tenants TN1 security alg sql
```

- Configure SUNRPC ALG.

```
[edit]
user@host# set tenants TN1 security alg sunrpc
```

- Configure TALK ALG.

```
[edit]
user@host# set tenants TN1 security alg talk
```

- Configure TFTP ALG.

```
[edit]
user@host# set tenants TN1 security alg tftp
```

- Configure TWAMP ALG.

```
[edit]
user@host# set tenants TN1 security alg twamp
```

- Configure extended function for FTP ALG.

```
[edit]
user@host# set tenants TN1 security alg ftp allow-mismatch-ip-address
```

- Configure extended function for MSRPC ALG.

```
[edit]
user@host# set tenants TN1 security alg msrpc map-entry-timeout 10
```

- Configure extended function for SUNRPC ALG.

```
[edit]
user@host# set tenants TN1 security alg sunrpc map-entry-timeout 10
```

- Configure extended function for SIP ALG.

```
[edit]
user@host# set tenants TN1 security alg sip retain-hold-resource
```

Example: Configuring ALG in Tenant System

IN THIS SECTION

- Requirements | 720
- Overview | 720
- Configuration | 720

- Verification | 725

This example shows how to configure ALGs in tenant system and send traffic based on FTP ALG configuration of the tenant system individually.

Requirements

This example uses the following hardware and software components:

- An SRX device
- Junos OS Release 18.3R1

Before you begin:

- Read the ALG Support for Tenant System to understand how and where this procedure fits in the overall tenant support for ALGs.

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, the ALG for FTP is configured to monitor and allow FTP traffic to be exchanged between the clients and the server on a tenant system.

By default, the FTP ALG is enabled on the tenant system.

Configuration

IN THIS SECTION

- CLI Quick Configuration | 721
- Configuring FTP ALG in a Tenant System | 721
- Results | 723

CLI Quick Configuration

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

```
set system security-profile p1 policy maximum 100
set system security-profile p1 policy reserved 50
set system security-profile p1 zone maximum 100
set system security-profile p1 zone reserved 50
set system security-profile p1 flow-session maximum 6291456
set system security-profile p1 flow-session reserved 50
set system security-profile p1 flow-gate maximum 524288
set system security-profile p1 flow-gate reserved 50
set tenants TN1 routing-instances VR_TN1 instance-type vpls
set tenants TN1 routing-instances VR_TN1 interface lt-0/0/0.0
set system security-profile p1 tenant TN1
set tenants TN1 security zones security-zone TN1_Czone host-inbound-traffic system-services all
set tenants TN1 security zones security-zone TN1_Czone host-inbound-traffic protocols all
set tenants TN1 security zones security-zone TN1_Czone interfaces ge-0/0/0
set tenants TN1 security zones security-zone TN1_Szone host-inbound-traffic system-services all
set tenants TN1 security zones security-zone TN1_Szone host-inbound-traffic protocols all
set tenants TN1 security zones security-zone TN1_Szone interfaces ge-0/0/1
set tenants TN1 security policies from-zone TN1_Czone to-zone TN1_Szone policy p11 match source-
address any
set tenants TN1 security policies from-zone TN1_Czone to-zone TN1_Szone policy p11 match
destination-address any
set tenants TN1 security policies from-zone TN1_Czone to-zone TN1_Szone policy p11 match
application junos-ftp
set tenants TN1 security policies from-zone TN1_Czone to-zone TN1_Szone policy p11 match
application junos-ping
set tenants TN1 security policies from-zone TN1_Czone to-zone TN1_Szone policy p11 then permit
set tenants TN1 security policies default-policy deny-all
```

Configuring FTP ALG in a Tenant System

Step-by-Step Procedure

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

To configure an ALG on a tenant system:

1. Configure a security profile p1 for tenant.

```
[edit]
set system security-profile p1 policy maximum 100
set system security-profile p1 policy reserved 50
set system security-profile p1 zone maximum 100
set system security-profile p1 zone reserved 50
set system security-profile p1 flow-session maximum 6291456
set system security-profile p1 flow-session reserved 50
set system security-profile p1 flow-gate maximum 524288
set system security-profile p1 flow-gate reserved 50
```

2. Configure interfaces and routing instances to the TN1.

```
[edit]
user@host# set tenants TN1 routing-instances VR_TN1 instance-type vpls
user@host# set tenants TN1 routing-instances VR_TN1 interface lt-0/0/0.0
```

3. Configure a security profile p1 and assign it to the tenant system TN1.

```
[edit]
user@host# set system security-profile p1 tenant TN1
```

4. Configure security zones and assign interfaces to each zone.

```
[edit]
user@host# set tenants TN1 security zones security-zone TN1_Czone host-inbound-traffic system-
services all
user@host# set tenants TN1 security zones security-zone TN1_Czone host-inbound-traffic
protocols all
user@host# set tenants TN1 security zones security-zone TN1_Czone interfaces ge-0/0/0
user@host# set tenants TN1 security zones security-zone TN1_Szone host-inbound-traffic system-
services all
user@host# set tenants TN1 security zones security-zone TN1_Szone host-inbound-traffic
protocols all
user@host# set tenants TN1 security zones security-zone TN1_Szone interfaces ge-0/0/1
```

5. Configure a security policy that permits FTP traffic from the TN1_Czone to-zone TN1_Szone.

```
[edit]
user@host# set tenants TN1 security policies from-zone TN1_Czone to-zone TN1_Szone policy p11
match source-address any
user@host# set tenants TN1 security policies from-zone TN1_Czone to-zone TN1_Szone policy p11
match destination-address any
user@host# set tenants TN1 security policies from-zone TN1_Czone to-zone TN1_Szone policy p11
match application junos-ftp
user@host# set tenants TN1 security policies from-zone TN1_Czone to-zone TN1_Szone policy p11
match application junos-ping
user@host# set tenants TN1 security policies from-zone TN1_Czone to-zone TN1_Szone policy p11
then permit
user@host# set tenants TN1 security policies default-policy deny-all
```

Results

From configuration mode, confirm your configuration by entering the `show tenants TN1` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show tenants TN1
    routing-instances {
        VR_TN1 {
            instance-type vpls;
            interface lt-0/0/0.0;
        }
    }
    security {
        policies {
            from-zone TN1_Czone to-zone TN1_Szone {
                policy p11 {
                    match {
                        source-address any;
                        destination-address any;
                        application [ junos-ftp junos-ping ];
                    }
                    then {
                        permit;
                    }
                }
```

```
            }
            default-policy {
                deny-all;
            }
        }
        zones {
            security-zone TN1_Czone {
                host-inbound-traffic {
                    system-services {
                        all;
                    }
                    protocols {
                        all;
                    }
                }
                interfaces {
                    ge-0/0/0.0;
                }
            }
            security-zone TN1_Szone {
                host-inbound-traffic {
                    system-services {
                        all;
                    }
                    protocols {
                        all;
                    }
                }
                interfaces {
                    ge-0/0/1.0;
                }
            }
        }
    }
}
```

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

Verification

IN THIS SECTION

- Verifying Intra-Tenant System traffic on ALG | 725
- Verify ALG status for Tenant System | 725

To confirm that the configuration is working properly, perform these tasks:

Verifying Intra-Tenant System traffic on ALG

Purpose

Verify the information about active resources, clients, groups, and sessions created through the resource manager.

Action

From operational mode, enter the `show security resource-manager summary` command.

```
user@host> show security resource-manager summary
Active resource-manager clients   : 0
Active resource-manager groups    : 0
Active resource-manager resources : 0
Active resource-manager sessions  : 0
```

Meaning

The output displays summary information about active resources, clients, groups, and sessions created through the resource manager.

Verify ALG status for Tenant System

Purpose

Verify the ALG status for tenant on the device.

Action

To verify the configuration is working properly, enter the `show security alg status tenant TN1` command.

```
user@host>show security alg status tenant TN1
ALG Status:
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Disabled
  SQL      : Disabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled
  TWAMP    : Disabled
```

Meaning

The output display the alg status for FTP Enabled for the tenant system TN1.

RELATED DOCUMENTATION

Tenant Systems Overview | 530

DHCP for Tenant Systems

IN THIS SECTION

- Understanding DHCP support for Tenant Systems | 727
- Minimum DHCPv6 Relay Agent Configuration for Tenant Systems | 727
- Example: Configuring a DHCPv6 Client for Tenant Systems | 729

Understanding DHCP support for Tenant Systems

Starting in Junos OS Release 18.4R1, a tenant system supports the DHCP client feature to learn IP addresses for interfaces assigned to the tenant systems. Additionally, starting in Junos OS Release 18.4R1, tenant systems support the DHCP relay feature. A DHCP relay agent forwards DHCP requests and responses between the DHCP client and the DHCP server.

An interface of an SRX Series device operating as a DHCP client receives the TCP or IP settings and the IP address from an external DHCP server.

An SRX Series device operating as a DHCP relay agent for tenant systems forwards incoming requests from the DHCP clients to a specified DHCP server. The client requests pass through interfaces on the tenant systems.

Minimum DHCPv6 Relay Agent Configuration for Tenant Systems

The following example describes the minimum configuration required to configure an SRX Series device as a DHCPv6 relay agent for the tenant system.

Before you begin determine the following:

- The DHCP routing instance name, the DHCP relay group and the DHCP active server-group for the tenant system.

1. Create a DHCPv6 relay group that includes at least one interface for the tenant system.

```
user@host# set tenants TSYS1 routing-instances R1 interface ge-0/0/0.0
```

2. Specify the DHCP group and add interfaces belonging to the group.

```
user@host# set tenants TSYS1  routing-instances R1 forwarding-options dhcp-relay dhcpv6 group
inf interface ge-0/0/0.0
```

3. Specify the name of the server-group and add the IP address for the DHCP servers belonging to the same group.

```
user@host# set tenants TSYS1  routing-instances R1 forwarding-options dhcp-relay dhcpv6
server-group server6 2001:db8::1/64
```

4. Specify the name of the active server-group.

```
user@host# set tenants TSYS1  routing-instances R1 forwarding-options dhcp-relay dhcpv6
active-server-group server6
```

5. Confirm your configuration by entering the `show tenants TSYS1 routing-instances R1` command.

```
[edit]
user@host# show tenants TSYS1 routing-instances R1
forwarding-options {
    dhcp-relay {
        dhcpv6 {
            group inf {
                interface ge-0/0/0.0;
            }
            server-group {
                server6 {
                    2001:db8::1/64;
                }
            }
            active-server-group server6;
        }
    }
}
```

Example: Configuring a DHCPv6 Client for Tenant Systems

IN THIS SECTION

- Requirements | 729
- Overview | 729
- Configuration | 730
- Verification | 734

This example shows how to configure a device as a DHCPv6 client for tenant systems.

Requirements

This example uses the following hardware and software components:

- An SRX Series device
- Junos OS Release 18.4R1

Before you begin:

- Read the Understanding DHCP support for Tenant Systems to understand how and where this procedure fits in the overall tenant systems support for DHCP.

  No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, a tenant system administrator configures an SRX Series device as a DHCPv6 client for a tenant system.

The DHCPv6 client for a tenant system includes the following features:

- Identity association for non-temporary addresses (IA_NA)
- Identity association for prefix delegation (IA_PD)
- Autoconfig or stateful mode
- DHCP unique identifier (DUID)

Configuration

IN THIS SECTION

- CLI Quick Configuration | 730
- Configuring DHCPv6 Client in a Tenant System | 730
- Procedure | 731
- Results | 732

CLI Quick Configuration

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

```
set tenants TSYS1 security zones security-zone trust host-inbound-traffic system-services all
set tenants TSYS1 security zones security-zone trust host-inbound-traffic protocols all
set tenants TSYS1 security zones security-zone trust interfaces ge-0/0/0.0
set tenants TSYS1 routing-instances r1 instance-type virtual-router
set tenants TSYS1 routing-instances r1 interface ge-0/0/0.0
set tenants TSYS1 interfaces ge-0/0/0 unit 0 family inet6 dhcpv6-client client-type autoconfig
set tenants TSYS1 interfaces ge-0/0/0 unit 0 family inet6 dhcpv6-client client-type stateful
set tenants TSYS1 interfaces ge-0/0/0 unit 0 family inet6 dhcpv6-client client-ia-type ia-na
set tenants TSYS1 interfaces ge-0/0/0 unit 0 family inet6 dhcpv6-client client-ia-type ia-pd
set tenants TSYS1 interfaces ge-0/0/0 unit 0 family inet6 dhcpv6-client client-identifier duid-
type duid-ll
set tenants TSYS1 interfaces ge-0/0/0 unit 0 family inet6 dhcpv6-client req-option dns-server
set protocols router-advertisement interface ge-0/0/0.0
```

Configuring DHCPv6 Client in a Tenant System

Procedure

Step-by-Step Procedure

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

1. Configure security zones to permit traffic for a tenant system.

```
[edit tenants TSYS1 security zones]
user@host# set security-zone trust host-inbound-traffic system-services all
user@host# set security-zone trust host-inbound-traffic protocols all
user@host# set security-zone trust interfaces ge-0/0/0.0
```

2. Create a routing instance and assign the routing instance type to a tenant system.

```
[edit tenants TSYS1]
user@host# set routing-instances r1 instance-type virtual-router
```

3. Specify the interface name for the routing instance.

```
[edit tenants TSYS1]
user@host# set routing-instances r1 interface ge-0/0/0.0
```

4. Configure the DHCPv6 client type. The client type can be `autoconfig` or `stateful` for a tenant system.

   - To enable DHCPv6 auto configuration mode, configure the client type as `autoconfig`.

```
[edit tenants TSYS1 interfaces ge-0/0/0 unit 0 family inet6 dhcpv6-client]
user@host# set client-type autoconfig
```

   - For stateful address assignment, configure the client type as `stateful`.

```
[edit tenants TSYS1 interfaces ge-0/0/0 unit 0 family inet6 dhcpv6-client]
user@host# set client-type stateful
```

5. Specify the identity association type.

- To configure identity association for nontemporary address (IA_NA) assignment, specify the `client-ia type` as `ia-na`.

```
[edit tenants TSYS1 interfaces ge-0/0/0 unit 0 family inet6 dhcpv6-client]
user@host# set client-ia-type ia-na
```

- To configure identity association for prefix delegation (IA_PD), specify the `client-ia-type` as `ia-pd`.

```
[edit tenants TSYS1 interfaces ge-0/0/0 unit 0 family inet6 dhcpv6-client]
user@host# set client-ia-type ia-pd
```

6. Configure the DHCPv6 client identifier by specifying the DHCP unique identifier (DUID) type for the tenant system. The following DUID type is supported:

   - Link Layer address (duid-ll)

```
[edit tenants TSYS1 interfaces ge-0/0/0 unit 0 family inet6 dhcpv6-client]
user@host# set client-identifier duid-type duid-ll
```

7. Specify the DHCPv6 client requested option as `dns-server` for the tenant system.

```
[edit tenants TSYS1 interfaces ge-0/0/0 unit 0 family inet6 dhcpv6-client]
user@host# set req-option dns-server
```

8. Configure the router advertisement.

```
[edit]
user@host# set protocols router-advertisement interface ge-0/0/0.0
```

Results

- From configuration mode, confirm your configuration by entering the `show tenants TSYS1` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show tenants TSYS1
interfaces {
```

```
    ge-0/0/0 {
        unit 0 {
            family inet6 {
                dhcpv6-client {
                    client-type stateful;
                    client-ia-type ia-na;
                    client-ia-type ia-pd;
                    client-identifier duid-type duid-ll;
                    req-option dns-server;
                }
            }
        }
    }
}
routing-instances {
    r1 {
        instance-type virtual-router;
        interface ge-0/0/0.0;
    }
}
security {
    zones {
        security-zone trust {
            host-inbound-traffic {
                system-services {
                    all;
                }
                protocols {
                    all;
                }
            }
            interfaces {
                ge-0/0/0.0;
            }
        }
    }
}
```

- From configuration mode, confirm your configuration by entering the `show protocols` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show protocols
router-advertisement {
    interface ge-0/0/0.0;
}
```

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

Verification

IN THIS SECTION

- Verifying the DHCPv6 Client for the Tenant System | 734
- Verifying the DHCPv6 Client Binding for the Tenant System | 735
- Verifying the DHCPv6 Client Statistics Information for the Tenant System | 736

To confirm that the configuration is working properly, perform these tasks:

Verifying the DHCPv6 Client for the Tenant System

Purpose

Verify that the DHCPv6 client information is configured.

Action

From the operational mode, enter the `show dhcpv6 client binding tenant TSYS1` command.

```
user@host> show dhcpv6 client binding  tenant TSYS1
IP/prefix      Expires     State      ClientType     Interface      Client DUID
         2000::17/128    67762    BOUND      STATEFUL         ge-0/0/6.0
LL0x3-10:0e:7e:49:25:86
```

```
          2000:100::/64   67762    BOUND      STATEFUL         ge-0/0/6.0
LL0x3-10:0e:7e:49:25:86
```

Meaning

The output displays the address binding information for the tenant system.

Verifying the DHCPv6 Client Binding for the Tenant System

Purpose

Verify that the DHCPv6 client binding information is configured.

Action

From the operational mode, enter the `show dhcpv6 client binding detail tenant TSYS1` command.

```
user@host> show dhcpv6 client binding detail tenant TSYS1
Client Interface/Id: ge-0/0/6.0
     Hardware Address:             10:0e:7e:49:25:86
     State:                        BOUND(DHCPV6_CLIENT_STATE_BOUND)
     ClientType                    STATEFUL
     Lease Expires:                2018-11-09 07:11:47 UTC
     Lease Expires in:             67760 seconds
     Lease Start:                  2018-11-08 07:11:47 UTC
     Bind Type:                    IA_NA IA_PD
     Preferred prefix length       0
     Sub prefix length             0
     Client DUID:                  LL0x3-10:0e:7e:49:25:86
     Rapid Commit:                 Off
     Server Identifier:            fe80::46f4:77ff:fed6:670a
     Client IP Address:            2000::17/128
     Client IP Prefix:             2000:100::/64

     DHCP options:
     Name: server-identifier, Value: VENDOR0x00000583-0x34343a34
```

Meaning

The output displays the detailed client binding information for the tenant system.

Verifying the DHCPv6 Client Statistics Information for the Tenant System

Purpose

Verify that the DHCP client statistics information is configured.

Action

From the operational mode, enter the `show dhcpv6 client statistics tenant TSYS1` command.

```
user@host> show dhcpv6 client statistics tenant TSYS1 routing-instance R1
Dhcpv6 Packets dropped:
    Total               3
    Bad Send            3

    Messages received:
    DHCPV6_ADVERTISE           1
    DHCPV6_REPLY               1
    DHCPV6_RECONFIGURE         0

    Messages sent:
    DHCPV6_DECLINE             0
    DHCPV6_SOLICIT             1
    DHCPV6_INFORMATION_REQUEST 0
    DHCPV6_RELEASE             0
    DHCPV6_REQUEST             1
    DHCPV6_CONFIRM             0
    DHCPV6_RENEW               0
    DHCPV6_REBIND              0
```

Meaning

The output displays the information about the number of packets discarded, the number of messages received and the number of messages sent by the DHCP client for the tenant system.

Security Log for Tenant Systems

IN THIS SECTION

- Understanding of Security Log for Tenant Systems | 737
- Example: Configure Security Log for Tenant Systems | 739
- Understanding On-Box Reporting for Tenant Systems | 744
- Configuring On-Box Reporting for Tenant Systems | 745
- Understanding On-Box and Off-Box Logging for Tenant System | 746
- Configuring On-Box Binary Security Log Files for Tenant System | 747
- Configuring Off-Box Binary Security Log Files for Tenant System | 748

Security logs for tenant systems include security events to control system's data planes. Security logs are sent in binary format to an external server from a tenant system interface. Security logs are generated per tenant system.

Understanding of Security Log for Tenant Systems

Junos OS generates separate log messages to record events that occur on the system's control and data planes. The data plane logs, also called security logs, primarily include security events that are handled inside the data plane. Security logs can be in text or binary format and they can be saved locally (event mode) or sent to an external server (stream mode). The binary format is required for stream mode and recommended to conserve log space in event mode.

If you configure security logs per tenant, then security logs are generated per tenant.

Security logs for a tenant system are sent from a tenant system interface. You can configure the assigned routing instances and the interfaces that belong to the routing tables within a tenant system.

A security profile should be defined with the number of maximum and reserved policies when you configure the stream number for a tenant system. The primary administrator can use the security profiles to specify resource allocation.

If a tenant system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available and not allocated to other tenant systems. The maximum allowed quota for stream number specifies the portion of the free global

resources that the tenant system can use. The maximum allowed quota does not ensure that the amount specified for the resource in the security profile is available. A reserved quota ensures that the resource amount specified is always available to the tenant system. Table 39 on page 738 shows the comparison of logging stream number capacity.

Table 39: Comparison of Logging Stream Number

Platform	Logging Stream Number Capacity for Tenant System + Logical System	Reserved Logging Stream Number Quota for Tenant System	Maximum Allowed Stream Number Quota for Tenant System	Maximum Allowed Stream Number Quota for Global
SRX5400, SRX5600, and SRX5800	64	0	8	64
SRX4600	300	0	8	600
SRX4100 and 4200	200	0	8	400
SRX1500	50	0	8	100

If a device is configured for a tenant system, security logs generated within the context have the **_LS** suffix in the log name, which is the same as the logical system. The following security log shows the attributes of the RT_FLOW_SESSION_CLOSE_LS log for a device that is configured for a tenant system:

```
<14>1 2018-03-12T22:50:09.596Z user RT_FLOW_SESSION_CLOSE_LS [junos@2636.1.1.1.2.137 logical-
system-name="TSYS1" reason="Some reason" source-address="192.0.2.1" source-port="7000"
destination-address="198.51.100.2" destination-port="32768" connection-tag="0" service-
name="Fake service" nat-source-address="192.0.2.1" nat-source-port="7000" nat-destination-
address="198.51.10    0.2" nat-destination-port="32768" nat-connection-tag="0" src-nat-rule-
type="Fake src nat rule" src-nat-rule-name="Fake src nat rule" dst-nat-rule-type="Fake dst nat
rule" dst-nat-rule-name="Fake dst nat rule" protocol-id="17" policy-name="Fake policy" source-
zone-name="Fake src zone" destination-zone-name="Fake dst zone" session-id-32="1" packets-from-
client="4294967295" bytes-from-client="4294967293" packets-from-server="4294967294" bytes-from-
server="4294967292" elapsed-time="4294967291" application="Fake application" nested-
application="Fake nested application" username="Fake username" roles="Fake UAC roles" packet-
incoming-interface="Fake packet incoming if" encrypted="Fake info telling if the traffic is
encrypted" application-category="Fake application category" application-sub-category="Fake
application subcategory" application-risk="-1"]
```

In the above example, security log includes **TSYS1** as the first attribute.

Starting in Junos OS Release 19.1R1, on-box reporting configurations are supported for each tenant system and logs are handled based on these configurations. Configure the `set security log report` and `set security log mode stream` commands to enable the on-box reporting. The on-box reporting feature with stream mode is also supported on tenant systems.

You can view Syslog messages in the System Log Explorer.

Example: Configure Security Log for Tenant Systems

IN THIS SECTION

- Requirements | **739**
- Overview | **739**
- Configuration | **740**
- Verification | **743**

This example shows how to configure security logs for a tenant system.

Requirements

This example uses the following hardware and software components:

- An SRX Series device.
- Junos OS Release 18.3R1 and later releases.

Before you begin:

- Understand how to configure a tenant system with security profiles for the primary logical system and two tenant systems. See Figure 1
-

Overview

SRX Series devices have two types of log: system logs and security logs. System logs record control plane events, for example, admin login to the device. Security logs, also known as traffic logs, record

data plane events regarding specific traffic handling, for example when a security policy denies certain traffic due to some violation of the policy.

The two types of logs can be collected and saved either on-box or off-box. The procedure below explains how to configure security logs in binary format for off-box (stream-mode) logging.

For off-box logging, security logs for a tenant system are sent from a tenant system interface. If the tenant system interface is already configured in a routing instance, then configure `routing-instance` *routing-instance-name* at `edit tenants` *tenant-name* `security log stream` *log-stream-name* `host` hierarchy. If the interface is not configured in routing instance, then no routing instance should be configured at `set tenants` *tenant-name* `security log stream` *log-stream-name* `host` hierarchy.

Configuration

IN THIS SECTION

- CLI Quick Configuration | **740**
- Procedure | **741**
- Procedure | **742**
- Results | **742**

CLI Quick Configuration

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

```
set tenants TSYS1 security log mode stream
set tenants TSYS1 security log stream TN1_s format binary host 1.3.54.22
set tenants TSYS1 security log source-address 2.3.45.66
set tenants TSYS1 security log transport protocol tls
set tenants TSYS1 routing-instances TN1_ri instance-type virtual-router
set tenants TSYS1 routing-instances TN1_ri interface ge-0/0/3
set tenants TSYS1 security log stream TN1_s host routing-instance TN1_ri
set system security-profile p1 security-log-stream-number reserved 1
set system security-profile p1 security-log-stream-number maximum 2
set system security-profile p1 tenant TSYS1
```

Procedure

Step-by-Step Procedure

The following procedure specifies how to configure security logs for a tenant system.

1. Specify the logging mode and the format for the log file. For off-box, stream-mode logging.

```
[edit ]
user@host# set tenants TSYS1 security log mode stream
user@host# set tenants TSYS1 security log stream TN1_s format binary host 1.3.54.22
```

2. For off-box security logging, specify the source address, which identifies the SRX Series device that generated the log messages. The source address is required.

```
[edit ]
user@host# set tenants TSYS1 security log source-address 2.3.45.66
```

3. Specify the routing instance and define the interface.

```
[edit ]
user@host# set tenants TSYS1 routing-instances TN1_ri instance-type virtual-router
user@host# set tenants TSYS1 routing-instances TN1_ri interface ge-0/0/3
```

4. Define routing instance for a tenant system. If the interface is already configured in routing instance, then configure `routing-instance` *`routing-instance-name`* at `edit tenants` *`tenant-name`* `security log stream` *`log-stream-name`* `host` hierarchy. If the interface is not configured in routing instance, then no routing instance should be configured at `set tenants` *`tenant-name`* `security log stream` *`log-stream-name`* `host` hierarchy.

```
[edit ]
user@host# set tenants TSYS1 security log stream TN1_s host routing-instance TN1_ri
```

5. Specify the security log transport protocol for the device.

```
[edit ]
user@host# set tenants TSYS1 security log transport protocol tls
```

Procedure

Step-by-Step Procedure

The following procedure specifies how to configure a security profile for a tenant system.

1. Configure a security profile and specify the number of maximum and reserved policies.

```
[edit ]
user@host# set system security-profile p1 security-log-stream-number reserved 1
user@host# set system security-profile p1 security-log-stream-number maximum 2
```

2. Assign the configured security profile to TSYS1.

```
[edit ]
user@host# set system security-profile p1 tenant TSYS1
```

Results

From configuration mode, confirm your configuration by entering the `show system security-profile`, `show tenants TSYS1 security log`, and `show tenants TSYS1 routing-instances` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show tenants TSYS1 security log
mode stream;
source-address 2.3.45.66;
transport {
    protocol tls;
}
stream TN1_s {
    format binary;
    host {
        1.3.54.22;
        routing-instance TN1_ri;
```

```
    }
}
```

```
[edit]
user@host# show tenants TSYS1 routing-instances
TN1_ri {
    instance-type virtual-router;
    interface ge-0/0/3.0;
}
```

```
[edit]
user@host# show system security-profile
p1 {
    security-log-stream-number {
        maximum 2;
        reserved 1;
    }
    tenant TSYS1;
}
```

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

Verification

IN THIS SECTION

- Verifying Detailed Output for Security Log | **743**

Verifying Detailed Output for Security Log

Purpose

Verify that the output displays the resource information for all tenant systems.

Action

From operational mode, enter the `show system security-profile security-log-stream-number tenant all` command.

```
logical-system tenant name  security profile name  usage    reserved    maximum

root-logical-system         Default-Profile           0          0         8
TSYS1                       p1                        1          1         2
```

Meaning

The output displays the resource information for tenant systems.

Understanding On-Box Reporting for Tenant Systems

Starting in Junos OS Release 19.1R1, on-box reporting configurations are supported for tenant systems and logs are handled based on these configurations.

Stream mode is a set of logging services that includes:

- Off-box logging (SRX Series)
- On-box logging and reporting (SRX1500, SRX4100, SRX4200, and SRX4600 Series)

Per tenant system configuration is supported for the off-box logging and logs are handled based on these configurations. The tenant system logs for off-box logging can only be generated from the tenant system interface.

On-box reporting mechanism is an enhancement to the existing logging functionality. The existing logging functionality is modified to collect system traffic logs, analyzes the logs, and generate reports of these logs. On-box reporting feature is intended to provide a simple and easy to use interface for viewing security logs.

Configure the `set security log report` and `set security log mode stream` commands to enable the on-box reporting feature on the device for tenant systems. The on-box reporting feature with stream mode is also supported on tenant systems.

The on-box reporting feature supports:

- Generating reports based on the requirements. For example: count or volume of the session, types of logs for activities such as IDP, UTM, and IPsec VPN.
- Capturing real-time events within a specified time range.
- Capturing all the network activities in a logical, organized, and easy-to-understand format based on various CLI specified conditions.

Configuring On-Box Reporting for Tenant Systems

SRX Series devices supports different types of reports for tenant system users.

Reports are stored locally on the SRX Series device and there is no requirement for separate devices or tools for logs and reports storage. The on-box reports provides a simple and easy-to-use interface for viewing the security logs.

Before you begin:

- Understand how to configure security log for tenant systems. See Example: Configure Security Log for Tenant Systems.

To configure on-box reporting for tenant system:

1. Define the tenant system name as TSYS1.

```
user@host# set tenants TSYS1
```

2. Create report within security log per tenant system.

```
user@host# set tenants TSYS1 security log report
```

3. Confirm your configuration by entering the `show tenants TSYS1` command.

```
user@host# show tenants TSYS1
security {
    log {
        report;
    }
}
```

NOTE: By default the `report` option is disabled.

Understanding On-Box and Off-Box Logging for Tenant System

SRX Series devices have two types of log: system logs and security logs. System logs record control plane events, for example admin login to the device. Security logs, also known as traffic logs, record data plane events regarding specific traffic handling, for example when a security policy denies certain traffic due to some violation of the policy.

Starting in Junos OS Release 19.2R1, on-box logging configurations are supported for each tenant system and logs are handled based on these configurations.

The two types of log can be collected and saved either on-box or off-box.

Stream mode is a set of logging services that includes:

- Off-box logging (SRX Series)
- On-box logging (SRX1500, SRX4100, SRX4200, and SRX4600 Series)

Per tenant system configuration is supported for the off-box logging and logs are handled based on these configurations. The tenant system logs for off-box logging can only be generated from the tenant system interface.

Configure the security files in **binary/syslog/sd-syslog/welf** format for stream-mode and binary format for event-mode by using the log statement at the `[set tenants TSYS1 security]` hierarchy level.

NOTE: You cannot configure the security log file path for Tenant System.

For on-box logging with stream mode with binary format log, the `set security log stream` *`stream-name`* `file` command is configured per tenant system. The file name must be end with **.bin**. For example **TSYS1_f1.bin** in tenant system TSYS1. A new file **TSYS1_f1.bin** is created in the **/var/traffic-log/tenant-systems/** ***TSYS1*** directory.

For on-box logging with stream mode with other format logs, the `set security log stream` *`stream-name`* `file` command is configured per tenant system. For example tenant system TSYS1. A new file with the name configured is created in the **/var/traffic-log/tenant-systems/** ***TSYS1*** directory.

Configuring On-Box Binary Security Log Files for Tenant System

SRX Series devices support two types of log: system logs and security logs.

The two types of log are collected and saved either on-box or off-box. The following procedure explains how to configure security logs in binary format for on-box (event-mode and stream-mode) logging for tenant system.

The following procedure specifies binary format for event-mode security logging, and defines the log filename, path, and log file characteristics for tenant system.

1. Specify the logging mode and the format for the log file. For on-box, event-mode logging:

```
[edit]
user@host# set tenants TSYS1 security log mode event
user@host# set tenants TSYS1 security log format binary
```

2. (Optional) Specify a log filename.

```
[edit]
user@host# set tenants TSYS1 security log file name security-binary-log
```

> **NOTE**: Security log filename is not mandatory. If security log filename is not configured, by default the file bin_messages is created in the /var/log directory.

3. Confirm your configuration by entering the `show tenants TSYS1` command.

```
[edit]
user@host# show tenants TSYS1
security {
    log {
        mode event;
        format binary;
        file {
            name security-binary-log;
        }
    }
}
```

The following procedure specifies binary format for stream-mode security logging, and defines the log filename and log file characteristics for tenant system.

1. Specify the logging mode and the format for the log file. For on-box, stream-mode logging:

```
[edit]
user@host# set tenants TSYS1 security log mode stream
user@host# set tenants TSYS1 security log stream s1 format binary
```

2. (Optional) Specify a log filename.

```
[edit]
user@host# set tenants TSYS1 security log stream s1 file name f1.bin
```

3. Confirm your configuration by entering the `show tenants TSYS1` command.

```
[edit]
user@host# show tenants TSYS1
security {
    log {
        mode stream;
        stream s1 {
            format binary;
            file {
                name f1.bin;
            }
        }
    }
}
```

Configuring Off-Box Binary Security Log Files for Tenant System

SRX Series devices support two types of log: system logs and security logs.

The two types of log can be collected and saved either on-box or off-box. The procedure below explains how to configure security logs in binary format for off-box (stream-mode) logging.

The following procedure specifies binary format for stream-mode security logging, and defines the logging mode, source address, and host name characteristics for tenant system.

1. Specify the logging mode and the format for the log file. For off-box, stream-mode logging:

```
[edit]
user@host# set tenants TSYS1 security log mode stream s1 format binary
```

2. Specify the source address for off-box security logging.

```
[edit]
user@host# set tenants TSYS1 security log source-address 100.0.0.1
```

3. Specify the host name.

```
[edit]
user@host# set tenants TSYS1 security log stream s1 host 100.0.0.2
```

4. Confirm your configuration by entering the `show tenants TSYS1` command.

```
[edit]
user@host# show tenants TSYS1
security {
    log {
        mode stream;
        source-address 100.0.0.1;
        stream s1 {
            format binary;
            host {
                100.0.0.2;
            }
        }
    }
}
```

Release History Table

Release	Description
19.2R1	Starting in Junos OS Release 19.2R1, on-box logging configurations are supported for each tenant system and logs are handled based on these configurations

AppQoS for Tenant Systems

IN THIS SECTION

- Application Quality of Service for Tenant Systems Overview | **750**
- Example: Configure Application Quality of Service for Tenant Systems | **751**

Application quality of service (AppQoS) enable you to identify and control access to specific applications and provides the granularity of the stateful firewall rule base to match and enforce quality of service (QoS) at the application layer. AppQoS feature expands the capability of Junos OS class of service (CoS) for tenant systems.

Application Quality of Service for Tenant Systems Overview

The application quality of service (AppQoS) feature expands the capability of Junos OS class of service (CoS) for tenant systems. This includes marking DSCP values based on Layer-7 application types, honoring application-based traffic through loss priority settings, and controlling transfer rates on egress PICs based on Layer-7 application types.

When a network experiences congestion and delay, some packets must be dropped. Junos OS CoS allows you to divide traffic into classes and offer various levels of throughput and packet loss when congestion occurs. This allows packet loss to happen according to the rules you configure.

Tenant system enables you to partition a single device into multiple domains to perform security and routing functions.

Starting in Junos OS Release 19.3R1, AppQoS is supported when the SRX Series device is configured with tenant system. You can configure a default AppQoS rule set to manage the application- traffic-

control within the tenant system. AppQoS provides the ability to prioritize and meter the application traffic to provide better service to business-critical or high-priority application traffic.

AppQoS rule sets are included in the tenant system to implement application-aware quality-of-service control. You can configure a rule set with rules under the application-traffic-control option, and attach the AppQoS rule set to a tenant system as an application service. If the traffic matches the specified application the application-aware quality of service is applied for tenant system.

For AppQoS, traffic is grouped based on rules that associate a defined forwarding class with selected applications for tenant system. The match criteria for the rule includes one or more applications. When traffic from a matching application encounters the rule, the rule action sets the forwarding class, and remarks the DSCP value and loss priority to values appropriate for the application.

The AppQoS DSCP rewriter conveys a packet's quality of service through both the forwarding class and a loss priority. The AppQoS rate-limiting parameters control the transmission speed and volume for its associated queues for tenant system. The default AppQoS rule set is leveraged from one of the existing AppQoS rule sets, which are configured under the `[edit class-of-service application-traffic-control]` hierarchy level.

Rate limiters are applied in rules based on the application of the traffic for tenant system. Two rate limiters are applied for each session: `client-to-server` and `server-to-client`. This usage allows traffic in each direction to be provisioned separately.

Example: Configure Application Quality of Service for Tenant Systems

IN THIS SECTION

- Requirements | **751**
- Overview | **752**
- Configuration | **752**
- Verification | **756**

This example shows how to enable application quality of service (AppQoS) within a tenant system to provide prioritization and rate limiting for the traffic.

Requirements

This example uses the following hardware and software components:

- An SRX Series device configured with tenant systems.
- Junos OS Release 19.3R1 and later releases.

Before you begin:

- Read the "Application Quality of Service for Tenant Systems Overview" on page 750 to understand how and where this procedure fits in the overall support for AppQos.

  No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you configure an AppQoS rule set and invoke AppQoS as an application service in the tenant systems. You configure the class of service (CoS) for tenant systems. The AppQoS rule sets are included in the tenant systems to implement application-aware quality-of-service control.

Configuration

IN THIS SECTION

- CLI Quick Configuration | **752**
- Configuring AppQoS with a Tenant System | **753**
- Results | **754**

CLI Quick Configuration

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

```
set tenants TSYS1 class-of-service application-traffic-control rate-limiters HTTP-BW-RL
bandwidth-limit 512
set tenants TSYS1 class-of-service application-traffic-control rule-sets RS1 rule RL1 match
application junos:HTTP
set tenants TSYS1 class-of-service application-traffic-control rule-sets RS1 rule RL1 then
forwarding-class best-effort
set tenants TSYS1 class-of-service application-traffic-control rule-sets RS1 rule RL1 then dscp-
code-point 001000
set tenants TSYS1 class-of-service application-traffic-control rule-sets RS1 rule RL1 then loss-
priority high
```

```
set tenants TSYS1 class-of-service application-traffic-control rule-sets RS1 rule RL1 then log
set tenants TSYS1 class-of-service application-traffic-control rule-sets RS1 rule RL1 then rate-
limit server-to-client HTTP-BW-RL
set tenants TSYS1 security policies from-zone untrust to-zone trust policy from_internet match
source-address any
set tenants TSYS1 security policies from-zone untrust to-zone trust policy from_internet match
destination-address any
set tenants TSYS1 security policies from-zone untrust to-zone trust policy from_internet match
application any
set tenants TSYS1 security policies from-zone trust to-zone trust policy p1 match dynamic-
application junos:web
set tenants TSYS1 security policies from-zone untrust to-zone trust policy from_internet then
permit application-services application-traffic-control rule-set RS1
```

Configuring AppQoS with a Tenant System

Step-by-Step Procedure

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

To configure AppQoS for a tenant system:

1. Configure the AppQoS real-time run information about application rate limiting of current or recent sessions for tenant system TSYS1.

```
user@host# set tenants TSYS1 class-of-service application-traffic-control rate-limiters HTTP-
BW-RL bandwidth-limit 512
```

2. Configure the AppQoS rules and application match criteria for tenant system TSYS1.

```
user@host# set tenants TSYS1 class-of-service application-traffic-control rule-sets RS1 rule
RL1 match application junos:HTTP
```

3. Configure the AppQoS rules and the forwarding class for tenant system TSYS1.

```
user@host# set tenants TSYS1 class-of-service application-traffic-control rule-sets RS1 rule
RL1 then forwarding-class best-effort
```

4. Configure the AppQoS rules and the dscp-code-point for tenant system TSYS1.

```
user@host# set tenants TSYS1 class-of-service application-traffic-control rule-sets RS1 rule
RL1 then dscp-code-point 001000
```

5. Configure the AppQoS rules and the loss priority for tenant system TSYS1.

```
user@host# set tenants TSYS1 class-of-service application-traffic-control rule-sets RS1 rule
RL1 then loss-priority high
```

6. Assign the rate limiters for rule-sets.

```
user@host# set tenants TSYS1 class-of-service application-traffic-control rule-sets RS1 rule
RL1 then log
user@host# set tenants TSYS1 class-of-service application-traffic-control rule-sets RS1 rule
RL1 then rate-limit server-to-client HTTP-BW-RL
```

7. Assign the class-of-service rule set to the security policy for tenant system TSYS1.

```
user@host# set tenants TSYS1 security policies from-zone untrust to-zone trust policy
from_internet match source-address any
user@host# set tenants TSYS1 security policies from-zone untrust to-zone trust policy
from_internet match destination-address any
user@host# set tenants TSYS1 security policies from-zone untrust to-zone trust policy
from_internet match application any
user@host# set tenants TSYS1 security policies from-zone trust to-zone trust policy p1 match
dynamic-application junos:web
user@host# set tenants TSYS1 security policies from-zone untrust to-zone trust policy
from_internet then permit application-services application-traffic-control rule-set RS1
```

Results

From configuration mode, confirm your configuration by entering the `show tenants TSYS1` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show tenants TSYS1
security {
```

```
    policies {
        from-zone untrust to-zone trust {
            policy from_internet {
                match {
                    source-address any;
                    destination-address any;
                    application any;
                }
                then {
                    permit {
                        application-services {
                            application-traffic-control {
                                rule-set RS1;
                            }
                        }
                    }
                }
            }
        }
        from-zone trust to-zone trust {
            policy p1 {
                match {
                    dynamic-application junos:web;
                }
            }
        }
    }
}
class-of-service {
    application-traffic-control {
        rate-limiters HTTP-BW-RL {
            bandwidth-limit 512;
        }
        rule-sets RS1 {
            rule RL1 {
                match {
                    application junos:HTTP;
                }
                then {
                    forwarding-class best-effort;
                    dscp-code-point 001000;
                    loss-priority high;
                    rate-limit {
```

```
                    server-to-client HTTP-BW-RL;
                }
                log;
            }
        }
    }
}
```

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

Verification

IN THIS SECTION

- Verifying the class-of-service application-traffic-control counter | 756
- Verifying the class-of-service application-traffic-control statistics rate-limiter | 757

To confirm that the configuration is working properly, perform the below tasks:

Verifying the class-of-service application-traffic-control counter

Purpose

Verify the class-of-service application-traffic-control counter for tenant systems.

Action

To verify the configuration is working properly, enter the `show class-of-service application-traffic-control counter tenant TSYS1` command.

```
user@host>show class-of-service application-traffic-control counter tenant TSYS1
Tenant System: TSYS1

pic: 0/0
  Counter type                                                  Value
 Sessions processed                                             1
 Sessions marked                                                0
```

```
Sessions honored                                    0
Sessions rate limited                               0
Client-to-server flows rate limited                 0
Server-to-client flows rate limited                 0
Session default ruleset hit                         0
Session ignored no default ruleset                  0
```

Meaning

The output displays AppQoS DSCP marking and honoring statistics based on Layer 7 application classifiers.

Verifying the class-of-service application-traffic-control statistics rate-limiter

Purpose

Verify the class-of-service application-traffic-control statistics rate-limiter for tenant systems.

Action

To verify the configuration is working properly, enter the `show class-of-service application-traffic-control statistics rate-limiter tenant TSYS1` command.

```
user@host>show class-of-service application-traffic-control statistics rate-limiter tenant TSYS1
Tenant System: TSYS1

pic: 0/0
```

Meaning

The output displays AppQoS real-time run information about application rate limiting of current or recent sessions.

Application Security for Tenant Systems

IN THIS SECTION

- Application Identification Services for Tenant Systems Overview | 758

Application Security in tenant systems identifies application traffic traversing your network regardless of port, protocol, and encryption, and thereby provides greater visibility to control network traffic. The application security controls network traffic by setting and enforcing security policies based on accurate application information.

Application Identification Services for Tenant Systems Overview

Predefined and custom application signatures identify an application by matching patterns in the first few packets of a session. Identifying applications provides the following advantages:

- Allows Intrusion Detection and Prevention (IDP) to apply appropriate attack objects to applications running on nonstandard ports.
- Improves performance by narrowing the scope of attack signatures for applications without decoders.
- Enables you to create detailed reports using AppTrack on applications passing through the device.

With tenant systems, predefined and custom application signatures are global resources that are shared by all tenant systems. Application identification (AppID) is enabled by default for a tenant system. The following are the privileges and responsibilities of the primary administrator over AppID:

- Download and install the predefined Juniper Networks application signatures.
- Create custom application and nested application signatures to identify applications that are not a part of the predefined database.
- View or clear the application identification statistics and counters for all tenant systems.
- Uninstall application signature package.

The application system cache (ASC) saves the mapping between an application type and the corresponding destination IP address, destination port, protocol type, and service. Each tenant system

has its own ASC. A tenant system or the primary administrator can view or clear ASC entries for any tenant system.

The AppID support for tenant systems include two options to view or clear tenant system statistics and tenant system counters for their own tenant system. Because the statistics reset time is common across the tenant systems, when you configure a new tenant system for the very first time, the statistics for that tenant system may get cleared even before the configured statistics reset time.

The custom signatures configured by the primary administrator can be configured in the tenant system security policies.

As a primary administrator or a tenant system user, you can use the commands `show services application-identification status` and `show services application-identification version` to view the status and version information about the AppID signature package.

4

CHAPTER

Configuration Statements

address-book (System) | 763

address-name | 765

anti-spam | 767

anti-virus | 769

auth-entry | 773

confidentiality | 775

content-filtering (Logical System Security Feature Profile) | 776

cpu | 779

dslite-softwire-initiator | 781

dynamic-address | 783

firewall-authentication (tenants) | 785

web-authentication | 787

pass-through | 789

flow-gate | 791

flow-session | 793

idp (logical-systems) | 796

idp-policy | 797

log (Security) | 799

log (Logical Systems and Tenant Systems) | 804

logical-system (System Security Profile) | 808

logical-domain-identity-management | 810

logical-systems (All) | 813

nat | 815

nat-cone-binding | 820

nat-destination-pool | 823

nat-destination-rule | 825

nat-interface-port-ol (System) | 827

nat-nopat-address | 829

nat-pat-address | 831

nat-pat-portnum | 833

nat-port-ol-ipnumber | 835

nat-rule-referenced-prefix (System) | 837

nat-source-pool | 839

nat-source-rule | 841

nat-static-rule | 843

policy (System Security Profile) | 845

policy-with-count | 848

profile | 850

protocols (Tenant Systems) | 851

purging | 853

root-authentication | 854

root-logical-system | 857

root-streaming | 858

secure-wire (System Security Profile) | 861

scheduler (System Security Profile) | 862

screen | 865

security-profile | 871

security-profile-resources | 876

stream (Logical Systems and Tenant Systems) | 878

softwires | 881

url | 883

web-filtering (Logical System Security Feature Profile) | 884

zone (System Security Profile) | 889

address-book (System)

IN THIS SECTION

- Syntax | 763
- Hierarchy Level | 763
- Description | 763
- Options | 764
- Required Privilege Level | 764
- Release Information | 764

Syntax

```
address-book {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level

```
[edit system security-profile security-profile-name]
```

Description

Address book entries include any combination of IPv4 addresses, IPv6 addresses, DNS names, wildcard addresses, and address range. You define addresses and address sets in an address book and then use those addresses when configuring different features, such as security policies and NAT.

Specify the number of address books that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

- `maximum` *amount*—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.
- `reserved` *amount*—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Address Books

address-name

IN THIS SECTION

- Syntax | 765
- Hierarchy Level | 766
- Description | 766
- Options | 766
- Required Privilege Level | 766
- Release Information | 767

Syntax

```
address-name name {
    description description;
    profile {
        category name {
            feed feed;
            property name {
                string name;
            }
        }
        feed-name name;
    }
}
```

Hierarchy Level

```
[edit security dynamic-address]
[edit logical-systems logical-system-name security dynamic-address]
[edit tenants tenant-name security dynamic-address]
```

Description

Specify the security dynamic address name for IPv4 and IPv6 networks within a logical system and tenant system.

Options

- `profile`—Information to categorize feed data into this dynamic address.
  - `feed-name`— Name of feed in feed-server for the dynamic address.
    - `source-feed-name`— Name of feed in feed-server which is mapped to this dynamic address.
  - `category`—Name of category for the dynamic address.
    - `GeoIP` —A list giving you the ability to filter traffic to and from specific geographies in the world.
      - `feed`—Feed to match.
        - `feed-name`—Name of feed (fd1).
    - `ip-filter` —A list of addresses and ranges of malicious sites that can send junk data.
      - `feed`—Feed to match.
        - `feed-name`—Name of feed (fd1).

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.4R1.

anti-spam

IN THIS SECTION

- Syntax | 767
- Hierarchy Level | 768
- Description | 768
- Options | 768
- Required Privilege Level | 769
- Release Information | 769

Syntax

```
anti-spam {
    sbl {
        profile (Security Antispam SBL) name {
            address-blacklist address-blacklist;
            address-whitelist address-whitelist;
            custom-tag-string custom-tag-string;
            (sbl-default-server | no-sbl-default-server);
            spam-action (block | tag-header | tag-subject);
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name security utm feature-profile]
[edit tenants tenant-name security utm feature-profile]
```

Description

Configures the UTM antispam feature for logical systems. The antispam feature examines transmitted e-mail messages to identify e-mail spam. When the device detects a message deemed to be spam, it blocks the e-mail message or tags the e-mail message header or subject with a preprogrammed string. Antispam filtering uses both a third-party server-based Spam Block List (SBL), and optionally created local allowlists (benign) and blocklists (malicious) for filtering against e-mail messages.

You can also configure the default UTM configuration for antispam feature profile. If you do not configure any option in the antispam feature profile, the values configured in the default UTM configuration are applied. In the default UTM profile, the antispam type is configured as SBL instead of none. This configuration enables the SBL option. However, to use this feature, enable the SBL server using the `[edit security utm default-configuration anti-spam sbl sbl-default-server]` command.

NOTE: A license check for the antispam configuration is performed at the time of commit and provides a warning if a valid license is not installed on the device. Once a valid license is installed on the device then the custom antispam profile or the default antispam profile is able to process the traffic. If a license is expired or is not installed, the antispam service does not process the traffic.

Options

anti-spam	Configures the UTM antispam feature for logical system.
address-blacklist	Enter an address blocklist custom object for local list spam filtering.
address-whitelist	Enter an address allowlist custom object for local list spam filtering.
sbl	Antispam filtering allows you to use both a third-party server-based SBL, and optionally created local allowlists and blocklists for filtering against e-mail messages.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

Support for configuration in tenant systems introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Antispam Filtering Overview

utm default-configuration

anti-virus

IN THIS SECTION

- Syntax | **770**
- Hierarchy Level | **771**
- Description | **771**
- Options | **771**
- Required Privilege Level | **772**
- Release Information | **772**

Syntax

```
anti-virus {
    profile name {
        fallback-options (Security Antivirus Sophos Engine) {
            content-size (Security Antivirus Sophos Engine) (block | log-and-permit | permit);
            decompress-err (block | log-and-permit | permit);
            default (Security Antivirus Sophos Engine) (block | log-and-permit | permit);
            engine-not-ready (Security Antivirus Sophos Engine) (block | log-and-permit | permit);
            out-of-resources (Security Antivirus Sophos Engine) (block | log-and-permit | permit);
            timeout (Security Antivirus Fallback Options Sophos Engine) (block | log-and-permit |
permit);
            too-many-requests (Security Antivirus Fallback Options Sophos Engine) (block | log-and-
permit | permit);
        }
        mime-whitelist {
            exception exception;
            list list;
        }
        notification-options (Security Antivirus) {
            fallback-block (Security Antivirus) {
                custom-message (Security Fallback Block) custom-message;
                custom-message-subject (Security Fallback Block) custom-message-subject;
                (notify-mail-sender (Security Fallback Block) | no-notify-mail-sender (Security Fallback
Block));
                type (Security Fallback Block) (message | protocol-only);
            }
            fallback-non-block (Security Antivirus) {
                custom-message (Security Fallback Non-Block) custom-message;
                custom-message-subject (Security Fallback Non-Block) custom-message-subject;
                (notify-mail-recipient | no-notify-mail-recipient);
            }
            virus-detection (Security Antivirus) {
                custom-message (Security Virus Detection) custom-message;
                custom-message-subject (Security Virus Detection) custom-message-subject;
                (notify-mail-sender (Security Virus Detection) | no-notify-mail-sender (Security Virus
Detection));
                type (Security Virus Detection) (message | protocol-only);
            }
        }
        url-whitelist url-whitelist;
```

```
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-systems-name security utm feature-profile]
[edit tenants tenant-name security utm feature-profile]
```

Description

Configures the UTM Sophos Antivirus feature for logical systems. You can also configure the default UTM configuration for antivirus feature profile. If you do not configure any option in the antivirus feature profile, the values configured in the default UTM configuration are applied.

NOTE: A license check for the antivirus configuration is performed at the time of a commit and will provide a warning if a valid license is not installed on the device. Once a valid license is installed on the device then a custom antivirus profile or the antivirus default profile is able to process traffic. If a license is expired or is not installed, the antivirus service does not process the traffic.

Options

anti-virus	Configures the UTM antivirus feature for logical systems.
mime-whitelist	This is the comprehensive list for those MIME types that can bypass antivirus scanning.
sophos-engine	The antivirus engine that is used on the device. You can only have one engine type running and you must restart the device if you change engines.
fallback-options	Fallback options helps the system how to handle the errors.

`notification-options`	There are multiple notification options you can configure to trigger when a virus is detected.
`fallback-non-block`	Notifications for fallback nonblocking actions.
`virus-detection`	Notifications to send when a virus is detected.
`scan-options`	Antivirus sophos-engine scan options.
`trickling`	HTTP trickling is a mechanism used to prevent the HTTP client or server from timing-out during a file transfer or during antivirus scanning.
`url-whitelist`	Antivirus URL allowlist is a unique custom list that includes the URLs or IP addresses category to bypass the antivirus scanning.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

Support for configuration in tenant systems introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Unified Threat Management Overview

utm default-configuration

auth-entry

IN THIS SECTION

- Syntax | 773
- Hierarchy Level | 773
- Description | 773
- Options | 774
- Required Privilege Level | 774
- Release Information | 774

Syntax

```
auth-entry {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the number of firewall authentication entries that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

- `maximum` *amount*—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.
- `reserved` *amount*—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Logical Systems Security Profiles (Primary Administrators Only) | 71

confidentiality

IN THIS SECTION

- Syntax | 775
- Hierarchy Level | 775
- Description | 775
- Required Privilege Level | 775
- Release Information | 776

Syntax

```
confidentiality;
```

Hierarchy Level

```
[edit security softwires map-e]
```

Description

Use this function to configure Junos OS MAP-E confidentiality.

Required Privilege Level

admin

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

show security softwires map-e confidentiality status | 1328

content-filtering (Logical System Security Feature Profile)

IN THIS SECTION

- Syntax | 776
- Hierarchy Level | 777
- Description | 777
- Options | 778
- Required Privilege Level | 778
- Release Information | 778

Syntax

```
content-filtering {
    profile name {
        block-command block-command;
        block-content-type {
            activex;
            exe;
            http-cookie;
            java-applet;
            zip;
```

```
        }
        block-extension block-extension;
        block-mime {
            exception (Security Content Filtering) exception;
            list (Security Content Filtering Block Mime) list;
        }
        notification-options (Security Content Filtering) {
            custom-message (Security Content Filtering) custom-message;
            (notify-mail-sender (Security Content Filtering Notification Options) | no-notify-mail-sender
(Security Content Filtering Notification Options));
            type (Security Content Filtering Notification Options) (message | protocol-only);
        }
        permit-command permit-command;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-systems-name security utm feature-profile]
[edit tenants tenant-name security utm feature-profile]
```

Description

Configures the UTM content-filtering feature for logical systems. The content filtering feature controls file transfers across the gateway by checking traffic against configured filter lists. It evaluates the traffic before all other UTM features, except Web filtering. You can also configure the default UTM configuration for content filtering feature profile. If you do not configure any option in the content filtering feature profile, the values configured in the default UTM configuration are applied.

NOTE: A license check for the content filtering configuration is performed at the time of commit and provides a warning if a valid license is not installed on the device. Once a valid license is installed on the device then the custom content filtering profile or the default content filtering profile is able to process the traffic. If a license is expired or license is not installed, the content filtering service does not process the traffic.

Options

`block-command`	Protocol block command custom-objects to the content-filtering profile.
`block-content-type`	Blocks other available content such as exe, http-cookie, java-applet. This is for HTTP only.
`block-extension`	Block extensions to the content-filtering profile.
`block-mime`	MIME pattern list custom-objects to the content-filtering profile for blocking MIME types.
`notification-options`	A message notification to trigger when a content filter is matched.
`permit-command`	Protocol permit command custom-objects to the content-filtering profile.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

Support for configuration in tenant systems introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Content Filtering Overview

utm default-configuration

cpu

IN THIS SECTION

- Syntax | 779
- Hierarchy Level | 779
- Description | 779
- Options | 780
- Required Privilege Level | 780
- Release Information | 780

Syntax

```
cpu {
    reserved percent;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the percentage of CPU utilization that is always available for a logical system. The available CPU space percentage value is configured in a security profile that is bound to a logical system.

Only the primary administrator can create security profiles and bind them to logical systems.

NOTE: The `cpu-control` option at the [`edit system security-profile resources`] hierarchy level must be specified for the `reserved` value to take effect.

Options

`reserved` *percent*—A reserved quota that guarantees that the percentage of CPU specified is always available to the logical system.

- **Range:** 0 through 100 percent (decimal point allowed).
- **Default:** 1 percent for the primary logical system and 0 percent for user logical systems.

CAUTION: The primary logical system must not be bound to a security profile that is configured with a 0 percent reserved CPU quota as traffic loss could occur.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Understanding Logical Systems Security Profiles (Primary Administrators Only) | 71

dslite-softwire-initiator

IN THIS SECTION

- Syntax | 781
- Hierarchy Level | 781
- Description | 781
- Options | 782
- Required Privilege Level | 782
- Release Information | 782

Syntax

```
dslite-softwire-initiator {
    maximum amount;
     reserved amount;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specifies the number of IPv6 dual-stack lite (DS-Lite) softwire initiators that connects to the softwire concentrator configured in either a user logical system or the primary logical system. The option is configured in the security profile that is bound to the logical system.

Only the primary administrator can create security profiles and bind them to logical systems. The primary administrator:

- Uses security profiles to provision logical systems with resources
- Binds security profiles to user logical systems and the primary logical system
- Configures more than one security profile, specifying different amounts of resource allocations in various profiles

Options

- `maximum` *amount*—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources. The default is the system maximum.
- `reserved` *amount*—A reserved quota that guarantees that the resource amount specified is always available to the logical system. The default is 0.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Understanding IPv6 Dual-Stack Lite in Logical Systems | 361

dynamic-address

IN THIS SECTION

- Syntax | 783
- Hierarchy Level | 784
- Description | 784
- Options | 784
- Required Privilege Level | 785
- Release Information | 785

Syntax

```
dynamic-address {
    address-name name {
        description description;
        profile {
            category name {
                feed feed;
                property name {
                    string name;
                }
            }
            feed-name name;
        }
    }
    feed-server feed-server-name {
         url url);
        description description;
        feed-name name {
            description description;
            hold-interval seconds;
            path path;
            update-interval seconds;
        }
```

```
        hold-interval seconds;
        hostname hostname;
        update-interval seconds;
    }
    traceoptions {
        file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
        flag name;
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit security]
[edit logical-systems logical-system-name security]
[edit tenants tenant-name security ]
```

Description

Configure security dynamic address for IPv4 and IPv6 networks within a logical system and tenant system. Each dynamic address belongs to only one instance. Within that instance is a set of categories to which the dynamic address further belongs. A dynamic address entry provides dynamic IP address information to the security policies.

Options

address-name	Security dynamic address name.
feed-server	Security dynamic address feed-server.
traceoptions	Security dynamic address tracing options.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.4R1.

firewall-authentication (tenants)

IN THIS SECTION

- Syntax | 785
- Hierarchy Level | 786
- Description | 787
- Required Privilege Level | 787
- Release Information | 787

Syntax

```
firewall-authentication {
    pass-through {
        default-profile default-profile;
        ftp {
            banner (Access FTP HTTP Telnet Authentication) {
                fail fail;
                login login;
                success success;
            }
```

```
        }
        ftp {
            banner (Access FTP HTTP Telnet Authentication) {
                fail fail;
                login login;
                success success;
            }
        }
        ftp {
            banner (Access FTP HTTP Telnet Authentication) {
                fail fail;
                login login;
                success success;
            }
        }
    }
    traceoptions {
        file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
        flag name;
        no-remote-trace;
    }
    web-authentication {
        banner {
            success success;
        }
        default-profile default-profile;
        timeout timeout;
    }
}
```

Hierarchy Level

```
[edit logical-systems name tenants name access],
[edit tenants name access]
```

Description

Define the type of firewall authentication on tenant system.

Required Privilege Level

access

Release Information

Statement introduced in Junos OS Release 18.3R1

RELATED DOCUMENTATION

- Firewall Authentication for Tenant Systems | 617
- show security firewall-authentication history | 1011
- show security firewall-authentication users | 1015
- firewall-authentication (tenants) | 785

web-authentication

IN THIS SECTION

- Syntax | 788
- Hierarchy Level | 788
- Description | 788
- Options | 788
- Required Privilege Level | 789
- Release Information | 789

Syntax

```
web-authentication {
    banner {
        success success;
    }
    default-profile default-profile;
    timeout timeout;
}
```

Hierarchy Level

```
[edit logical-systems name tenants name access firewall-authentication],
[edit tenants name access firewall-authentication]
```

Description

Define Web-authentication settings for tenant systems.

Options

default-profile Name of profile to use for web-authentication

timeout Web-authentication timeout value in seconds

- **Range:** 5 through 60

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

flow-tap

Release Information

Statement introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

pass-through | 789

pass-through

IN THIS SECTION

- Syntax | 789
- Hierarchy Level | 790
- Description | 790
- Options | 790
- Required Privilege Level | 791
- Release Information | 791

Syntax

```
pass-through {
    default-profile default-profile;
    ftp {
        banner (Access FTP HTTP Telnet Authentication) {
            fail fail;
```

```
            login login;
            success success;
        }
    }
    ftp {
        banner (Access FTP HTTP Telnet Authentication) {
            fail fail;
            login login;
            success success;
        }
    }
    ftp {
        banner (Access FTP HTTP Telnet Authentication) {
            fail fail;
            login login;
            success success;
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems name tenants name access firewall-authentication],
[edit tenants name access firewall-authentication]
```

Description

Pass-through firewall authentication settings

Options

default-profile Name of profile to use if not specified in policy

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

flow-tap

Release Information

Statement introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

firewall-authentication (tenants) | 785

flow-gate

IN THIS SECTION

- Syntax | 791
- Hierarchy Level | 792
- Description | 792
- Options | 792
- Required Privilege Level | 793
- Release Information | 793

Syntax

```
flow-gate {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specifies the number of flow gates, also known as `pinholes`, that user logical system administrators and primary logical system administrators can configure for their logical systems, if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

maximum *amount*
: A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **Range:** 0 through 1000000

reserved *amount*
: A reserved quota that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Logical Systems Security Profiles (Primary Administrators Only) | 71

flow-session

IN THIS SECTION

- Syntax | 793
- Hierarchy Level | 794
- Description | 794
- Options | 794
- Required Privilege Level | 795
- Release Information | 795

Syntax

```
flow-session {
    maximum amount;
```

```
    reserved amount;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specifies the number of flow sessions that user logical system administrators and primary logical system administrators configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

maximum *amount*
: A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

  - **Range:** 0 through 1000000

reserved *amount*
: A reserved quota that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

> **NOTE**: An IPv6 session consumes twice the memory of an IPv4 session. Therefore the number of sessions available for IPv6 is half the reserved and maximum quotas configured for the flow session resource in a security profile. Use the vty command **show usp flow resource usage cp-session** to check flow session usage.

For performance consideration, flow session rate-limiting is done on distributed central point (DCP) for SRX5400, SRX5600, and SRX5800 devices. DCP session is always created before SPU flow session on these devices. You can count DCP session instead of SPU flow session to check whether the maximum/reserved quotas are reached for a specific logical system or tenant system.

There are no differences between DCP session and SPU flow session count because a session without NAT consumes one DCP session and one SPU flow session.

For sessions with NAT configuration, one connection can consume one SPU flow session and two DCP sessions. If all sessions are NAT sessions and each session consumes two DCP sessions, then the logical system or tenant system can only reach half of the maximum quotas configured for the flow session resources.

Typically, there are sessions with NAT and non-NAT mixed traffic, so the logical system or tenant system can create flow sessions between “maximum”/2 and “maximum”.

This functionality is similar to the reserved quotas that are configured for the flow session resources in a security profile. As a logical system or tenant system can consume more resources than the reserved quotas, you can configure additional flow session resources for reserved quotas with NAT configuration, if there are free resources available in the logical system or tenant system.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Logical Systems Security Profiles (Primary Administrators Only) | 71

Example: Configuring Logical Systems Security Profiles (Primary Administrators Only) | 77

idp (logical-systems)

IN THIS SECTION

- Syntax | 796
- Hierarchy Level | 796
- Description | 796
- Options | 797
- Required Privilege Level | 797
- Release Information | 797

Syntax

```
idp (default | off | on);
```

Hierarchy Level

```
[edit logical-systems name security]
```

Description

Configure IDP on primary and user logical systems.

Options

default-policy	Set the active policy.
max-sessions	Maximum number of IDP sessions.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security

Release Information

Statement introduced in Junos OS Release 11.4

RELATED DOCUMENTATION

- show security idp logical system | **1037**
- Understanding IDP Features in Logical Systems | **276**
- IDP in Logical Systems Overview | **274**

idp-policy

IN THIS SECTION

- Syntax | **798**
- Hierarchy Level | **798**
- Description | **798**
- Options | **798**

Required Privilege Level | 798
Release Information | 799

Syntax

```
idp-policy idp-policy;
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the IDP policy for the security profile.

Options

idp-policy-name– Name of the IDP policy.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Intrusion Detection and Prevention Overview

log (Security)

IN THIS SECTION

- Syntax | 799
- Hierarchy Level | 802
- Description | 802
- Options | 802
- Required Privilege Level | 804
- Release Information | 804

Syntax

```
log {
    (source-address source-address | source-interface source-interface);
    cache {
        exclude name {
            destination-address destination-address;
            destination-port destination-port;
            event-id event-id;
            failure;
            interface-name interface-name;
            policy-name policy-name;
            process process;
```

```
            protocol protocol;
            source-address source-address;
            source-port source-port;
            success;
            username username;
        }
        limit limit;
    }
    disable;
    escape;
    time-format (year | millisecond);
    event-rate logs per second;
    facility-override (authorization | daemon | ftp | kernel | local0 | local1 | local2 | local3 |
local4 | local5 | local6 | local7 | user);
    file {
        files files;
        name name;
        path path;
        size size;
    }
    format (binary | sd-syslog | syslog);
    max-database-record max-database-record;
    message-rate-limit messages per second;
    mode (event | stream | stream-event);
    rate-cap logs per second;
    report {
        logs-per-table {
            idp idp;
            ipsec-vpn ipsec-vpn;
            screen screen;
            session-all session-all;
            sky sky;
            utm utm;
        }
        table-lifetime table-lifetime;
        table-mode {
            dense;
        }
    }
    root-streaming;
    stream stream-name {
        category (all | content-security | fw-auth | screen | alg | nat | flow | sctp | gtp |
ipsec | idp | rtlog |pst-ds-lite | appqos |secintel |aamw);
```

```
        filter {
            threat-attack;
            }
        format (binary | sd-syslog | syslog | welf);
        host {
            ip-address;
            port port-number;
            routing-instanceinstance-name;
        }
        rate-limit {
            log-rate;
        }
        severity (alert | critical | debug | emergency | error | info | notice | warning);
        source-address {
            ip-address;
        }
        time-format (year | millisecond);
        transport {
            protocol (tcp | tls | udp);
            tcp-connections tcp-connections;
            tls-profile tls-profile;
        }
    }
    traceoptions {
        file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
        flag name;
        no-remote-trace;
    }
    transport {
        protocol (tcp | tls | udp);
        tcp-connections tcp-connections;
        tls-profile tls-profile;
    }
    utc-timestamp;
}
```

Hierarchy Level

```
[edit security]
[edit logical-systems name security]
[edit tenants tenant-name security]
```

Description

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

Options

cache Cache security log events in the audit log buffer.

disable Disable the security logging for the device.

escape Escapes the stream log forwarding to avoid parsing errors. Stream mode supports escape in `sd-syslog` and `binary` format. Event mode supports escape only in `binary` format.

time-format Specify the year, the millisecond, or both in the timestamp.

event-rate *rate* Limit the rate at which logs are streamed per second.

- **Range:** 0 through 1500
- **Default:** 1500

facility-override Alternate facility for logging to remote host.

file Specify the security log file options for logs in binary format.

- Values:

- *max-file-number*—Maximum number of binary log files.
  - The range is 2 through 10 and the default value is 10.
- *file-name*—Name of binary log file.
- *binary-log-file-path*—Path to binary log files.
- *maximum-file-size*—Maximum size of binary log file in megabytes.
  - The range is 1 through 10 and the default value is 10.

format Set the security log format for the device.

max-database-record The following are the disk usage range limits for the database:

- Range:
  - SRX1500, SRX4100, and SRX4200: 0 through 15,000,000
  - vSRX: 0 through 1,000,000
- Default:
  - SRX1500, SRX4100, and SRX4200: 15,000,000
  - vSRX: 1,000,000

  Be sure there is enough free space in `/var/log/hostlogs/`, otherwise logs might be dropped when written into the database.

mode Control how security logs are processed and exported.

rate-cap ***rate-cap-value*** Work with event mode only. This option limits the rate at which data plane logs are generated per second.

- **Range:** 0 through 5000 logs per second
- **Default:** 5000 logs per second

root-streaming Allows the user logical systems to generate the logs using the root logical system's stream configuration.

source-address ***source-address*** Specify a source IP address or IP address used when exporting security logs, which is mandatory to configure *stream host*.

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

	The `source-address` and `source-interface` are alternate values. Using one of the options is mandatory.
`stream`	Every stream can configure file or host.
`traceoptions`	Specify security log daemon trace options.
`transport`	Set security log transport settings.
`utc-timestamp`	Specify to use UTC time for security log timestamps.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

The [edit `logical-systems` *name* security] and [edit `tenants` *tenant-name* security] hierarchy levels introduced in Junos OS Release 19.1R1.

`escape` option added in Junos OS Release 20.2R1.

`root-streaming` option added in Junos OS Release 20.3R1.

log (Logical Systems and Tenant Systems)

IN THIS SECTION

Syntax | 805

- Hierarchy Level | 806
- Description | 806
- Options | 806
- Required Privilege Level | 808
- Release Information | 808

Syntax

```
        log {
            disable;
            facility-override (authorization | daemon | ftp | kernel | local0 | local1 | local2
| local3 | local4 | local5 | local6 | local7 | user);
            mode (event | stream);
            format (binary | sd-syslog | syslog);
            source-address;
            source-interface;
            stream name {
                category name;
                filter name;
                host {
                    ipaddr;
                    port port;
                    routing-instance routing-instance;
                }
                rate-limit rate;
            }
            transport {
                protocol (tcp | tls |udp);
                tls-profile;
            }
        }
    }
```

Hierarchy Level

```
[edit logical-systems name security]
[edit tenants name security]
```

Description

Configure security log for logical systems and tenant systems. Set the mode of logging (event for traditional system logging or stream for streaming security logs through a revenue port to a server). You can also specify all the other parameters for security logging.

Options

disable Disable the security logging for the device.

facility-override Alternate facility for logging to remote host.

- Values:
  - `authorization` —Authorization system
  - `daemon` —Various system processes
  - `ftp` —FTP process
  - `kernel` —Kernel
  - `local0` —Local logging option number 0
  - `local1` —Local logging option number 1
  - `local2` —Local logging option number 2
  - `local3` —Local logging option number 3
  - `local4` —Local logging option number 4

- `local5` —Local logging option number 5
- `local6` —Local logging option number 6
- `local7` —Local logging option number 7
- `user` —User processes

format Set security log format for the device.

- Values:
  - `binary` —Binary log
  - `binarysd-syslog` —Structured syslog
  - `syslog` —Traditional syslog

mode Controls how security logs are processed and exported.

- Values:
  - `event` —Process security logs in the control plane
  - `stream` —Process security logs directly in the forwarding plane

source-address Specify a source IP address or IP address used when exporting security logs, which is mandatory to configure *stream host*.

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

stream Set security log stream settings.

transport Set security log transport settings.

- Values:
  - tcp—TCP transfer for log
  - tls—TLS transfer for log
  - udp—UDP transfer for log

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

The following options are not supported under logical system and tenant system:

- `event-rate` and `rate-cap`— Use to limit the log rate between Packet Forwarding Engine (PFE) and Routing Engine (RE).

- `file`— Use to store binary log with event mode.
- `max-database-record` and `report`— Use to enable SQLite Version 3 (sqlite3) database for local log management daemon (llmd).
- `traceoptions`— Specify security log daemon trace options.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

Statement introduced in Junos OS Release 18.3R1 for tenant systems.

The `routing-instance` option introduced in Junos OS Release 18.3R1 for tenant systems.

RELATED DOCUMENTATION

Understanding Security Logs and Logical Systems | 509

logical-system (System Security Profile)

IN THIS SECTION

- Syntax | 809
- Hierarchy Level | 809
- Description | 809

- Options | 809
- Required Privilege Level | 810
- Release Information | 810

Syntax

```
logical-system logical-system;
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the user logical system to bind the security profile.

The primary administrator uses security profiles to provision logical systems with resources. You can bind security profiles to user logical systems and the primary logical system. The primary administrator can configure more than one security profile allocating different amounts of a resource in various ones.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

logical-system-name–Name of the logical system.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Logical Systems for SRX Series Services Gateways | 5

logical-domain-identity-management

IN THIS SECTION

- Syntax | 810
- Hierarchy Level | 812
- Description | 812
- Options | 812
- Required Privilege Level | 812
- Release Information | 812

Syntax

```
logical-domain-identity-management {
    active {
        authentication-entry-timeout minutes;
```

```
        filter {
            domain name;
            exclude-ip {
                address-book book-name;
                address-set address-set;
            }
            include-ip {
                address-book book-name;
                address-set address-set;
            }
        }
        invalid-authentication-entry-timeout minutes;
        ip-query {
            query-delay-time seconds;
        }
        query-server name {
            batch-query {
                items-per-batch items-per-batch;
                query-interval seconds;
            }
            connection (Identity Management Advanced Query){
                connect-method (http | https);
                port port;
                primary {
                    address address;
                    ca-certificate ca-certificate;
                    client-id client-id;
                    client-secret client-secret;
                }
                query-api query-api;
                secondary {
                    address address;
                    ca-certificate ca-certificate;
                    client-id client-id;
                    client-secret client-secret;
                }
                token-api token-api;
            }
        }
    }
    traceoptions {
        file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
```

```
        flag name;
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Specify the logical domain identity management configuration for both logical systems and tenant systems.

Options

active Displays the active mode for logical domain identity management module.

traceoptions Displays the tracing options.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Understanding Integrated User Firewall support in a Logical System | 193

Understanding Integrated User Firewall Support in a Tenant System | 635

logical-systems (All)

IN THIS SECTION

- Syntax | 813
- Hierarchy Level | 813
- Description | 814
- Options | 814
- Required Privilege Level | 814
- Release Information | 814

Syntax

```
logical-systems {
    logical-system-name {
        ...logical-system-configuration...
    }
}
```

Hierarchy Level

```
[edit]
```

Description

Configure a logical system. Only the primary administrator can configure a logical system at [edit] hierarchy level.

You can include several of the hierarchies that can be included at the [edit] hierarchy level. For descriptions of the applicable statements, see the appropriate hierarchies.

NOTE: The logical-systems configuration statement can be used only by the primary administrator.

Options

logical-system-name—Name of the logical system.

Required Privilege Level

all—To view this statement in the configuration.

all—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Logical Systems for SRX Series Services Gateways | 5

nat

IN THIS SECTION

- Syntax | 815
- Hierarchy Level | 819
- Description | 819
- Options | 819
- Required Privilege Level | 820
- Release Information | 820

Syntax

```
nat {
    destination {
        pool pool-name {
            address ip-address {
                (port port-number | to ip-address);
            }
            description text;
            routing-instance routing-instance-name;
        }
        rule-set rule-set-name {
            description text;
            from {
                interface [interface-name];
                routing-instance [routing-instance-name];
                zone [zone-name];
            }
            rule rule-name {
                description text;
                match {
                    (destination-address <ip-address> | destination-address-name <address-name>);
                    destination-port port-number;
                    protocol [protocol-name-or-number];
```

```
                    source-address [ip-address];
                    source-address-name [address-name];
                }
                then {
                    destination-nat (off | pool pool-name);
                }
            }
        }
    }
    proxy-arp {
        interface interface-name {
            address ip-address {
                to ip-address;
            }
        }
    }
    proxy-ndp {
        interface interface-name {
            address ip-address {
                to ip-address;
            }
        }
    }
    source {
        address-persistent;
        interface {
            port-overloading {
                off;
            }
        }
        pool pool-name {
            address ip-address {
                to ip-address;
            }
            description text;
            host-address-base ip-address;
            overflow-pool (interface | pool-name);
            port {
                (no-translation | port-overloading-factor number | range port-low <to port-
high>);
            }
            routing-instance routing-instance-name;
        }
```

```
        pool-default-port-range lower-port-range to upper-port-range;
        pool-utilization-alarm {
            clear-threshold value;
            raise-threshold value;
        }
        port-randomization {
            disable;
        }
        port-round-robin {
            disable;
        }
        rule-set rule-set-name {
            description text;
            from {
                interface [interface-name];
                routing-instance [routing-instance-name];
                zone [zone-name];
            }
            rule rule-name {
                description text;
                match {
                    (destination-address <ip-address> | destination-address-name <address-name>);
                    destination-port port-number;
                    protocol [protocol-name-or-number];
                    source-address [ip-address];
                    source-address-name [address-name];
                }
                then {
                    source-nat {
                        interface {
                            persistent-nat {
                                address-mapping;
                                inactivity-timeout seconds;
                                max-session-number value;
                                permit (any-remote-host | target-host | target-host-port);
                            }
                        }
                        off;
                        pool {
                             persistent-nat {
                                 address-mapping;
                                inactivity-timeout seconds;
                                max-session-number number;
```

```
                            permit (any-remote-host | target-host | target-host-port);
                        }
                        pool-name;
                    }
                }
            }
        }
        to {
            interface [interface-name];
            routing-instance [routing-instance-name];
            zone [zone-name];
        }
    }
}
static {
    rule-set rule-set-name {
        description text;
        from {
            interface [interface-name];
            routing-instance [routing-instance-name];
            zone [zone-name];
        }
        rule rule-name {
            description text;
            match {
                (destination-address ip-address | destination-address-name address-name);
            }
            then {
                static-nat {
                    inet {
                        routing-instance (default | routing-instance-name);
                    }
                    prefix {
                        address-prefix;
                        routing-instance (default | routing-instance-name);
                    }
                    prefix-name {
                        address-prefix-name;
                        routing-instance (default | routing-instance-name);
                    }
                }
            }
        }
```

```
        }
    }
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag;
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit security]
[edit tenants tenant-name security]
```

Description

Configure Network Address Translation (NAT) for NFX Series and SRX Series devices.

Options

destination	Configure Destination NAT.
natv6v4	Configure NAT between IPv6 and IPv4 options.
proxy-arp	Configure Proxy ARP.
proxy-ndp	Configure Proxy NDP.

source	Configure Source NAT.
static	Configure Static NAT.
traceoptions	Configure NAT traceoptions.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.6.

The `description` option added in Junos OS Release 12.1.

The tenant option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Introduction to NAT

Understanding Logical Systems Network Address Translation | 133

nat-cone-binding

IN THIS SECTION

- Syntax | 821
- Hierarchy Level | 821
- Description | 821

Options | 822
Required Privilege Level | 822
Release Information | 822

Syntax

```
nat-cone-binding {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the number of NAT cone binding configurations that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

maximum A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **Range:** 0 through 2097152

reserved A reserved quota that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Introduction to NAT

nat-destination-pool

IN THIS SECTION

- Syntax | 823
- Hierarchy Level | 823
- Description | 823
- Options | 824
- Required Privilege Level | 824
- Release Information | 824

Syntax

```
nat-destination-pool {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile]
[edit tenant tenant-name security-profile]
```

Description

Specify the number of NAT destination pool configurations that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

maximum A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **Range:** 0 through 8192

reserved A reserved quota that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

The tenant option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Introduction to NAT

nat-destination-rule

IN THIS SECTION

- Syntax | 825
- Hierarchy Level | 825
- Description | 826
- Options | 826
- Required Privilege Level | 826
- Release Information | 827

Syntax

```
nat-destination-rule {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile]
[edit tenant tenant-name system security-profile]
```

Description

Specify the number of NAT destination rule configurations that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

maximum A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **Range:** 0 through 8192

reserved A reserved quota that guarantees that the resource amount specified is always available.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

The tenant option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Introduction to NAT

nat-interface-port-ol (System)

IN THIS SECTION

- Syntax | 827
- Hierarchy Level | 828
- Description | 828
- Options | 828
- Required Privilege Level | 828
- Release Information | 828

Syntax

```
nat-interface-port-ol {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the security NAT interface port overloading the quota of a logical system.

Options

maximum *amount* Specify the maximum allowed quota value.

- **Range:** 0 through 64

reserved *amount* Specify a reserved quota value that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Introduction to NAT

nat-nopat-address

IN THIS SECTION

- Syntax | 829
- Hierarchy Level | 829
- Description | 829
- Options | 830
- Required Privilege Level | 830
- Release Information | 830

Syntax

```
nat-nopat-address {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the number of NAT without port address translation configurations that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

maximum A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **Range:** 0 through 4194304

reserved A reserved quota that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Introduction to NAT

nat-pat-address

IN THIS SECTION

- Syntax | 831
- Hierarchy Level | 831
- Description | 832
- Options | 832
- Required Privilege Level | 832
- Release Information | 833

Syntax

```
nat-pat-address {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the number of NAT with port address translation (PAT) configurations that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

maximum A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **Range:** 0 through 2048

reserved A reserved quota that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Logical Systems Network Address Translation | 133

Introduction to NAT

nat-pat-portnum

IN THIS SECTION

- Syntax | 833
- Hierarchy Level | 834
- Description | 834
- Options | 834
- Required Privilege Level | 834
- Release Information | 835

Syntax

```
nat-pat-portnum {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the maximum quantity and the reserved quantity of ports for the logical system as part of its security profile. The total number of PAT pools must not exceed the configured maximum ports for the logical system.

Options

maximum Specify the maximum number of ports allowed for a logical system. The maximum quantity is not guaranteed and is shared among multiple logical systems.

- **Range:** 0 through 134217728

reserved Specify the number of resources guaranteed for a logical system.

- **Range:** For SRX5600 and SRX5800 devices, up to 402,653,184 ports are supported. Pool specifications for logical systems can be viewed using the **show security nat source summary logical-system all** command.
- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Understanding Logical Systems for SRX Series Services Gateways | 5

nat-port-ol-ipnumber

IN THIS SECTION

- Syntax | 835
- Hierarchy Level | 836
- Description | 836
- Options | 836
- Required Privilege Level | 837
- Release Information | 837

Syntax

```
nat-port-ol-ipnumber {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the number of NAT port overloading IP number configurations that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

maximum A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **Range:** 0 through 16

reserved A reserved quota that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Logical Systems for SRX Series Services Gateways | 5

nat-rule-referenced-prefix (System)

IN THIS SECTION

- Syntax | 837
- Hierarchy Level | 838
- Description | 838
- Options | 838
- Required Privilege Level | 838
- Release Information | 838

Syntax

```
nat-rule-referenced-prefix {
    maximum maximum;
```

```
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the security NAT rule referenced IP prefix quota of a logical system.

Options

maximum Specify the maximum allowed quota value.

- **Range:** 0 through 1048576

reserved Specify a reserved quota value that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Understanding Logical Systems for SRX Series Services Gateways | 5

nat-source-pool

IN THIS SECTION

- Syntax | 839
- Hierarchy Level | 839
- Description | 840
- Options | 840
- Required Privilege Level | 840
- Release Information | 841

Syntax

```
nat-source-pool {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile]
[edit tenant tenant-name system security-profile]
```

Description

Specify the NAT source pool configurations that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

maximum A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **Range:** 0 through 2048

reserved A reserved quota that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

The tenant option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Understanding Logical Systems for SRX Series Services Gateways | 5

nat-source-rule

IN THIS SECTION

- Syntax | 841
- Hierarchy Level | 842
- Description | 842
- Options | 842
- Required Privilege Level | 843
- Release Information | 843

Syntax

```
nat-source-rule {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile]
[edit tenant tenant-name system security-profile]
```

Description

Specify the NAT source rule configurations that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

maximum A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **Range:** 0 through 8192

reserved A reserved quota that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

The tenant option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Understanding Logical Systems for SRX Series Services Gateways | 5

nat-static-rule

IN THIS SECTION

- Syntax | 844
- Hierarchy Level | 844
- Description | 844
- Options | 844
- Required Privilege Level | 845
- Release Information | 845

Syntax

```
nat-static-rule {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile]
[edit tenant tenant-name system security-profile]
```

Description

Specify the number of NAT static rule configurations that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

maximum A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed

quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **Range:** 0 through 8192

reserved A reserved quota that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

The tenant option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Understanding Logical Systems for SRX Series Services Gateways | 5

policy (System Security Profile)

IN THIS SECTION

Syntax | 846
Hierarchy Level | 846

- Description | 846
- Options | 847
- Required Privilege Level | 847
- Release Information | 847

Syntax

```
policy {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the number of security policies that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

maximum A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **Range:** 0 through 80000

reserved A reserved quota that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Logical Systems for SRX Series Services Gateways | 5

policy-with-count

IN THIS SECTION

- Syntax | 848
- Hierarchy Level | 848
- Description | 848
- Options | 849
- Required Privilege Level | 849
- Release Information | 849

Syntax

```
policy-with-count {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the number of security policies with a count that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

maximum A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **Range:** 0 through 1024

reserved A reserved quota that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Logical Systems for SRX Series Services Gateways | 5

profile

IN THIS SECTION

- Syntax | 850
- Hierarchy Level | 850
- Description | 850
- Options | 851
- Required Privilege Level | 851
- Release Information | 851

Syntax

```
profile profile-name;
```

Hierarchy Level

```
[edit access]
```

Description

Configure an access profile that specifies authentication and accounting parameters. Authentication and accounting parameters run only if you specify the profile name. The access profile contains the entire

authentication, authorization, and accounting (AAA) configuration that aids in handling AAA requests, authentication method and order, AAA server addresses, and AAA accounting.

Options

profile-name Name of the access profile.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.2R1 at tenant system level.

RELATED DOCUMENTATION

Firewall Authentication for Tenant Systems | 617

protocols (Tenant Systems)

IN THIS SECTION

- Syntax | 852
- Hierarchy Level | 852
- Description | 852
- Options | 852

Required Privilege Level | 853
Release Information | 853

Syntax

```
protocols (bgp | ospf | ospf3);
```

Hierarchy Level

```
edit tenants tenant-name routing-instances instance_name
```

Description

Specify the types of routing protocol traffic that can reach a tenant system supported device.

Options

Tenant system supports the following routing protocols:

bgp Enable incoming BGP traffic.

ospf Enable incoming OSPF traffic.

ospf3 Enable incoming OSPF version 3 traffic.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Tenant Systems Overview | 530

purging

IN THIS SECTION

- Syntax | 853
- Hierarchy Level | 854
- Description | 854
- Required Privilege Level | 854
- Release Information | 854

Syntax

```
purging;
```

Hierarchy Level

```
[edit system arp]
```

Description

Purge obsolete ARP entries from the cache when an interface or link goes offline.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

root-authentication

IN THIS SECTION

- Syntax | 855
- Hierarchy Level | 855
- Description | 855
- Options | 856
- Required Privilege Level | 856
- Release Information | 856

Syntax

```
root-authentication {
    encrypted-password "password";
    no-public-keys
    ssh-ecdsa name {
        from from;
    }
    ssh-ed25519 name {
        from from;
    }
    ssh-rsa name {
        from from;
    }
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure the authentication methods for the root-level user, whose username is `root`.

You can use the `ssh-ecdsa`, `ssh-ed25519`, or `ssh-rsa` statements to directly configure SSH ECDSA, ED25519, or RSA keys to authenticate root logins. You can configure more than one public key for SSH authentication of root logins as well as for user accounts. When a user logs in as root, the public keys are referenced to determine whether the private key matches any of them.

Options

encrypted-password "*password*"	Specify the MD5 or other password. You can specify only one encrypted password. You cannot configure a blank password using blank quotation marks (" "). You must configure a password whose number of characters range from 1 through 128 characters and enclose the password in quotation marks.
no-public-keys	Disable SSH public key-based authentication.
ssh-ecdsa *name* from *from*	Use an SSH ECDSA public key. You can specify one or more public keys.
ssh-ed25519 *name* from *from*	Use an SSH ED25519 public key. You can specify one or more public keys.
ssh-rsa *name* from *from*	Use an SSH RSA public key. You can specify one or more public keys.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

class (Defining Login Classes)

user (Access)

root-logical-system

IN THIS SECTION

- Syntax | 857
- Hierarchy Level | 857
- Description | 857
- Options | 858
- Required Privilege Level | 858
- Release Information | 858

Syntax

```
root-logical-system;
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify root-logical-system to bind the security profile to the primary logical system.

The primary administrator uses security profiles to provision logical systems with resources. The security profile containing this statement must be bound to root-logical-system only.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

none

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Logical Systems for SRX Series Services Gateways | 5

root-streaming

IN THIS SECTION

- Syntax | 859
- Hierarchy Level | 859
- Description | 859
- Required Privilege Level | 860
- Release Information | 860

Syntax

```
root-streaming;
```

Hierarchy Level

```
[edit security log,
[edit logical-systems name security log],
[edit tenants name security log]
```

Description

The root-streaming option allows the user logical systems and tenant systems to generate the logs using the root system's stream configuration. If you configure root-streaming option under user logical systems and tenant systems in stream mode, all the logs of this user logical system and tenant systems will be sent out according to root logical and tenant system's configuration.

The root-streaming option works together with the stream mode. If user logical system is in event mode, the root-streaming option will not take effect. If user logical system is in stream mode and root logical systems is in event mode, the root-streaming option will try to get the streams in root but the logs will not be sent out. Table 1 provides a detailed behavior of user logical system logs.

Table 40: root-streaming option for root and logical systems in different mode

Logical systems mode	Root logical systems mode	Status of the root-streaming option for logical systems	Behavior of logical systems logs
Event mode		Enabled	Send to Routing Engine
Stream mode	Event mode	Enabled	Cannot send out
	Stream mode	Enabled	Send to root log server

Table 40: root-streaming option for root and logical systems in different mode *(Continued)*

Logical systems mode	Root logical systems mode	Status of the root-streaming option for logical systems	Behavior of logical systems logs
	Stream-event mode	Enabled	Send to root log server
Stream-event mode	Event mode	Enabled	Send to Routing Engine
	Stream mode	Enabled	Send to Routing Engine and root log server
	Stream-event mode	Enabled	Send to Routing Engine and root log server
Event mode		Not Enabled	Send to Routing Engine
Stream mode		Not Enabled	Send to logical systems log server
Stream-event mode		Not Enabled	Send to Routing Engine and logical systems log server

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.3R1.

secure-wire (System Security Profile)

IN THIS SECTION

- Syntax | 861
- Hierarchy Level | 861
- Description | 861
- Options | 862
- Required Privilege Level | 862
- Release Information | 862

Syntax

```
secure-wire {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify maximum and reserved quota of a user logical system. If the user logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if the resources are available—that is, if they are not allocated to other user logical systems.

Options

maximum A maximum allowed quota of a user logical system.

- **Range:** 0 through 255

reserved A reserved quota of a user logical system.

- **Default:** 0

Required Privilege Level

admin

Release Information

Statement introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Secure Wire for Logical Systems | 227

scheduler (System Security Profile)

IN THIS SECTION

- Syntax | 863
- Hierarchy Level | 863
- Description | 863
- Options | 864
- Required Privilege Level | 864

Release Information | 864

Syntax

```
scheduler {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the number of schedulers that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

maximum A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **Range:** 0 through 64

reserved A reserved quota that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Logical Systems for SRX Series Services Gateways | 5

screen

IN THIS SECTION

- Syntax | 865
- Hierarchy Level | 869
- Description | 870
- Options | 870
- Required Privilege Level | 870
- Release Information | 870

Syntax

```
screen {
    ids-option name {
        aggregation {
            destination-prefix-mask destination-prefix-mask;
            destination-prefix-v6-mask destination-prefix-v6-mask;
            source-prefix-mask source-prefix-mask;
            source-prefix-v6-mask source-prefix-v6-mask;
        }
        alarm-without-drop;
        description (Security Screen) description;
        icmp (Security Screen) {
            flood (Security ICMP) <threshold ICMP packets per second>;
            fragment;
            icmpv6-malformed;
            ip-sweep <threshold microseconds in which 10 ICMP packets are detected>;
            large;
            ping-death;
        }
        ip (Security Screen) {
            bad-option;
            block-frag;
            ipv6-extension-header {
```

```
                AH-header;
                destination-header {
                    home-address-option;
                    ILNP-nonce-option;
                    line-identification-option;
                    tunnel-encapsulation-limit-option;
                    user-defined-option-type name {
                        to type-high;
                    }
                }
                ESP-header;
                fragment-header;
                HIP-header;
                hop-by-hop-header {
                    CALIPSO-option;
                    jumbo-payload-option;
                    quick-start-option;
                    router-alert-option;
                    RPL-option;
                    SMF-DPD-option;
                    user-defined-option-type name {
                        to type-high;
                    }
                }
                mobility-header;
                no-next-header;
                routing-header;
                shim6-header;
                user-defined-header-type name {
                    to type-high;
                }
            }
            ipv6-extension-header-limit ipv6-extension-header-limit;
            ipv6-malformed-header;
            loose-source-route-option;
            record-route-option;
            security-option;
            source-route-option;
            spoofing;
            stream-option;
            strict-source-route-option;
            tear-drop;
            timestamp-option;
```

```
            tunnel (Security Screen) {
                bad-inner-header;
                gre {
                    gre-4in4;
                    gre-4in6;
                    gre-6in4;
                    gre-6in6;
                }
                ip-in-udp {
                    teredo;
                }
                ipip {
                    dslite;
                    ipip-4in4;
                    ipip-4in6;
                    ipip-6in4;
                    ipip-6in6;
                    ipip-6over4;
                    ipip-6to4relay;
                    isatap;
                }
            }
            unknown-protocol;
        }
        limit-session {
            by-destination {
                by-protocol {
                    icmp {
                        maximum-sessions maximum-sessions;
                        packet-rate packet-rate;
                        session-rate session-rate;
                    }
                    tcp {
                        maximum-sessions maximum-sessions;
                        packet-rate packet-rate;
                        session-rate session-rate;
                    }
                    udp {
                        maximum-sessions maximum-sessions;
                        packet-rate packet-rate;
                        session-rate session-rate;
                    }
                }
```

```
                maximum-sessions maximum-sessions;
                packet-rate packet-rate;
                session-rate session-rate;
            }
            by-source {
                by-protocol {
                    icmp {
                        maximum-sessions maximum-sessions;
                        packet-rate packet-rate;
                        session-rate session-rate;
                    }
                    tcp {
                        maximum-sessions maximum-sessions;
                        packet-rate packet-rate;
                        session-rate session-rate;
                    }
                    udp {
                        maximum-sessions maximum-sessions;
                        packet-rate packet-rate;
                        session-rate session-rate;
                    }
                }
                maximum-sessions maximum-sessions;
                packet-rate packet-rate;
                session-rate session-rate;
            }
            destination-ip-based destination-ip-based;
            source-ip-based source-ip-based;
        }
        match-direction (input | input-output | output);
        tcp (Security Screen) {
            fin-no-ack;
            land;
            port-scan <threshold microseconds in which 10 attack packets are detected>;
            syn-ack-ack-proxy <threshold un-authenticated connections>;
            syn-fin;
            syn-flood {
                alarm-threshold requests per second;
                attack-threshold proxied requests per second;
                destination-threshold SYN pps;
                source-threshold SYN pps;
                timeout (Security Screen) seconds;
                white-list name {
```

```
                    destination-address [ destination-address ... ];
                    source-address [ source-address ... ];
                }
            }
            syn-frag;
            tcp-no-flag;
            tcp-sweep <threshold microseconds in which 10 TCP packets are detected>;
            winnuke;
        }
        udp (Security Screen) {
            flood (Security UDP) {
                threshold UDP packets per second;
                white-list [ white-list ... ];
            }
            port-scan <threshold microseconds in which 10 attack packets are detected>;
            udp-sweep <threshold microseconds in which 10 UDP packets are detected>;
        }
    }
    traceoptions (Security Screen) {
        file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
        flag name;
        no-remote-trace;
    }
    trap <interval seconds>;
    white-list name {
        address [ address ... ];
    }
}
```

Hierarchy Level

```
[edit security]
[edit tenant tenant-name security]
```

Description

Configure the security screen options. For every security zone, you can enable a set of predefined screen options that detect and block various kinds of traffic that the device determines as potentially harmful.

Options

ids-options *screen-name* — Name of the screen configured at the `security screen ids-options` level. Define screens for the intrusion detection service (IDS).

trap — Configure trap interval. Enable or disable the sending of Simple Network Management Protocol (SNMP) notifications when the state of the connection changes. Traps are unsolicited messages sent from an SNMP agent to remote network management systems or trap receivers.

white-list — Set of IP addresses for allowlist. Configure an allowlist of IP addresses that are to be exempt from the SYN cookie and SYN proxy mechanisms that occur during the SYN flood screen protection process. An allowlist contains known trusted IP addresses and URLs. Content downloaded from locations on the allowlist does not have to be inspected for malware.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

The `description` option added in Junos OS Release 12.1.

The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

security-profile

IN THIS SECTION

- Syntax | 871
- Hierarchy Level | 874
- Description | 874
- Options | 875
- Required Privilege Level | 875
- Release Information | 875

Syntax

```
security-profile {
    profile name {
        address-book (System) {
            maximum maximum;
            reserved reserved;
        }
        advanced-anti-malware-policy {
            maximum maximum;
            reserved reserved;
        }
        auth-entry {
            maximum maximum;
            reserved reserved;
        }
        cpu {
            reserved percent;
        }
```

```
        dslite-softwire-initiator {
            maximum maximum;
            reserved reserved;
        }
        flow-gate {
            maximum maximum;
            reserved reserved;
        }
        flow-session {
            maximum maximum;
            reserved reserved;
        }
        icap-redirect-profile {
            maximum maximum;
            reserved reserved;
        }
        idp-policy idp-policy;
        logical-system (System Security Profile) logical-system;
        nat-cone-binding {
            maximum maximum;
            reserved reserved;
        }
        nat-destination-pool {
            maximum maximum;
            reserved reserved;
        }
        nat-destination-rule {
            maximum maximum;
            reserved reserved;
        }
        nat-interface-port-ol (System) {
            maximum maximum;
            reserved reserved;
        }
        nat-nopat-address {
            maximum maximum;
            reserved reserved;
        }
        nat-pat-address {
            maximum maximum;
            reserved reserved;
        }
        nat-pat-portnum {
```

```
            maximum maximum;
            reserved reserved;
        }
        nat-port-ol-ipnumber {
            maximum maximum;
            reserved reserved;
        }
        nat-rule-referenced-prefix (System) {
            maximum maximum;
            reserved reserved;
        }
        nat-source-pool {
            maximum maximum;
            reserved reserved;
        }
        nat-source-rule {
            maximum maximum;
            reserved reserved;
        }
        nat-static-rule {
            maximum maximum;
            reserved reserved;
        }
        policy (System Security Profile) {
            maximum maximum;
            reserved reserved;
        }
        policy-with-count {
            maximum maximum;
            reserved reserved;
        }
        root-logical-system;
        scheduler (System Security Profile) {
            maximum maximum;
            reserved reserved;
        }
        secintel-policy {
            maximum maximum;
            reserved reserved;
        }
        secure-wire {
            maximum maximum;
            reserved reserved;
```

```
        }
        security-log-stream-number {
            maximum maximum;
            reserved reserved;
        }
        tenant tenant;
        user-auth-entry {
            maximum maximum;
            reserved reserved;
        }
        vrf-group {
            maximum maximum;
            reserved reserved;
        }
        zone (System Security Profile) {
            maximum maximum;
            reserved reserved;
        }
    }
    resources {
        cpu-control;
        cpu-control-target percent;
    }
}
```

Hierarchy Level

```
[edit system]
[edit tenants <tenant-name>]
```

Description

Create a security profile and specify the kinds and amounts of resources to allocate to a logical system to which the security profile is bound.

As a primary administrator, you can create a security profile and bind it to more than one logical system if you want to allocate the same kinds and amounts of resources to them. For details on how many

security profiles you can create, see "Understanding Logical Systems Security Profiles (Primary Administrators Only)" on page 71. When you reach the limit, you must delete a security profile and commit the configuration before you can create and commit the configuration for another security profile.

Only the primary administrator can create security profiles.

Options

profile-name Name of the security profile.

resources Name of the resources.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

The `dslite-softwire-initiator` option introduced in Junos OS Release 12.1.

The `security-profile` option added under the `tenants` hierarchy in Junos OS Release 18.3R1.

The icap redirect profile option is introduced in Junos OS Release 18.3R1.

`secure-wire` option introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Understanding Logical Systems Security Profiles (Primary Administrators Only) | 71

Example: Configuring Logical Systems Security Profiles (Primary Administrators Only) | 77

security-profile-resources

IN THIS SECTION

- Syntax | 876
- Hierarchy Level | 876
- Description | 876
- Options | 877
- Required Privilege Level | 877
- Release Information | 877

Syntax

```
security-profile-resources {
    cpu-control;
    cpu-control-target percent;
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure global settings that apply to all logical systems in the device.

Options

`cpu-control`—Enable CPU utilization control.

`cpu-control-target` *percent*—Specify the upper limit for CPU utilization on the device under normal operating conditions.

- **Range:** 0 through 100 percent (decimal point allowed).
- **Default:** 80 percent.

> **NOTE**: The `cpu-control` option must be specified for the `cpu-control-target` value to take effect.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Understanding Logical Systems Security Profiles (Primary Administrators Only) | 71

Example: Configuring Logical Systems Security Profiles (Primary Administrators Only) | 77

stream (Logical Systems and Tenant Systems)

IN THIS SECTION

- Syntax | 878
- Hierarchy Level | 878
- Description | 879
- Options | 879
- Required Privilege Level | 880
- Release Information | 880

Syntax

```
stream name {
    category name;
    filter name;
    host {
        ipaddr;
        port port;
        routing-instance routing-instance;
    }
    rate-limit rate;
}
```

Hierarchy Level

```
[edit logical-systems name security log]
[edit tenants name security log]
```

Description

Define security log stream options for a logical and tenant system. When the logging mode is set to `stream`, security logs generated in the data plane are streamed out a revenue traffic port directly to a remote server. All the categories can be configured for sending specific category logs to different log servers for stream mode log forwarding.

Options

`category` Type of logged events.

- Values:
  - aamw—AAMW events are logged.
  - alg—Alg events are logged.
  - all—All events are logged.
  - appqos—Appqos events are logged.
  - content-security—Content security events are logged.
  - flow—Flow events are logged.
  - fw-auth—Fw-auth events are logged.
  - gtp—Gtp events are logged.
  - idp—Idp events are logged.
  - ipsec—Ipsec events are logged.
  - nat—Nat events are logged.
  - pst-ds-lite—Pst-ds-lite events are logged.
  - rtlog—Rtlog events are logged.
  - screen—Screen events are logged.
  - sctp—Sctp events are logged.
  - secintel—Secintel events are logged.

filter Selects the filter to filter the logs to be logged.

- Values:
  - `threat-attack` — Threat attack security events are logged.

host Destination to send security logs.

- *ip-address*—Specify IP address of the host.
- `port` *port-number*—Specify host port number.
- `routing-instance` *routing-instance*—Specify the routing instance.

rate-limit Rate-limit for security logs.

- **Range:** 1 through 65535
- **Default:** 65535

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1 for logical Systems.

Statement introduced in Junos OS Release 18.3R1.

The `routing-instance` option introduced in Junos OS Release 18.3R1 for tenant systems.

RELATED DOCUMENTATION

Understanding Security Logs and Logical Systems | 509

softwires

IN THIS SECTION

- Syntax | 881
- Hierarchy Level | 882
- Description | 882
- Options | 882
- Required Privilege Level | 882
- Release Information | 883

Syntax

```
softwires {
    rule-set name {
        match-direction (input | output);
        rule name {
            then {
                v6rd v6rd;
            }
        }
    }
    softwire-name name {
        ipv4-prefix ipv4-prefix;
        mtu-v4 mtu-v4;
        softwire-concentrator softwire-concentrator;
        softwire-type (IPv4-in-IPv6 | v6rd);
        v6rd-prefix v6rd-prefix;
    }
    traceoptions {
        file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
        flag name;
        no-remote-trace;
```

```
    }
}
```

Hierarchy Level

```
[edit logical-systems name security],
[edit security]
[edit services]
```

Description

Configure softwires for IPv6 dual-stack lite (DS-Lite). DS-Lite allows migration to an IPv6 access network without changing end-user software. IPv4 users can continue to access IPv4 internet content using their current hardware, while IPv6 users are able to access IPv6 content.

Options

map-e	Name of the map-e rules configuration.
softwire-name	Name of the softwire configuration.
traceoptions	Traceoptions for Network Security DS-Lite.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 12.1.

url

IN THIS SECTION

- Syntax | 883
- Hierarchy Level | 883
- Description | 883
- Required Privilege Level | 884
- Release Information | 884

Syntax

```
url url;
```

Hierarchy Level

```
[edit security dynamic-address feed-server feed-server-name]
```

Description

URL of the bundle file for feed server configuration. Use a browser to verify the URL validity. Examples of URLs are shown below:

- example.url.com/my_file.tgz
- example.url.com/my_file.bundle

Required Privilege Level

security

Release Information

Statement introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Dynamic Address Groups in Security Policies

web-filtering (Logical System Security Feature Profile)

IN THIS SECTION

- Syntax | **885**
- Hierarchy Level | **887**
- Description | **887**
- Options | **887**
- Required Privilege Level | **889**
- Release Information | **889**

Syntax

```
web-filtering {
    juniper-enhanced {
        profile name {
            base-filter base-filter;
            block-message {
                type custom-redirect-url;
                url url;
            }
            category name {
                action (block | log-and-permit | permit | quarantine);
                custom-message custom-message;
            }
            custom-block-message custom-block-message;
            default (block | log-and-permit | permit | quarantine);
            fallback-settings {
                default (block | log-and-permit);
                server-connectivity (block | log-and-permit);
                timeout (block | log-and-permit);
                too-many-requests (block | log-and-permit);
            }
            no-safe-search;
            quarantine-custom-message quarantine-custom-message;
            quarantine-message {
                type custom-redirect-url;
                url url;
            }
            site-reputation-action {
                fairly-safe (block | log-and-permit | permit | quarantine);
                harmful (block | log-and-permit | permit | quarantine);
                moderately-safe (block | log-and-permit | permit | quarantine);
                suspicious (block | log-and-permit | permit | quarantine);
                very-safe (block | log-and-permit | permit | quarantine);
            }
            timeout seconds;
        }
    }
    juniper-local {
        profile name {
            block-message {
```

```
                type custom-redirect-url;
                url url;
            }
            category name {
                action (block | log-and-permit | permit | quarantine);
                custom-message custom-message;
            }
            custom-block-message custom-block-message;
            default (block | log-and-permit | permit);
            fallback-settings {
                default (block | log-and-permit);
                server-connectivity (block | log-and-permit);
                timeout (block | log-and-permit);
                too-many-requests (block | log-and-permit);
            }
            quarantine-custom-message quarantine-custom-message;
            quarantine-message {
                type custom-redirect-url;
                url url;
            }
            timeout seconds;
        }
    }
    websense-redirect {
        profile name {
            account account;
            block-message {
                type custom-redirect-url;
                url url;
            }
            category name {
                action (block | log-and-permit | permit | quarantine);
                custom-message custom-message;
            }
            custom-block-message custom-block-message;
            fallback-settings {
                default (block | log-and-permit);
                server-connectivity (block | log-and-permit);
                timeout (block | log-and-permit);
                too-many-requests (block | log-and-permit);
            }
            quarantine-custom-message quarantine-custom-message;
            quarantine-message {
```

```
                type custom-redirect-url;
                url url;
            }
            server {
                host host;
                port port;
                routing-instance routing-instance;
            }
            sockets sockets;
            timeout seconds;
        }
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-systems-name security utm feature-profile]
[edit tenants tenant-name security utm feature-profile]
```

Description

Configures the UTM Web filtering feature for logical systems. The Web filtering allows you to manage Internet usage by preventing access to inappropriate Web content. The potential policies conflict check of the Web filtering feature is independent of the content filtering, antivirus, and antispam features. You can also configure the default UTM configuration for Web filtering feature profile. If you do not configure any option in the Web filtering feature profile, the values configured in the default UTM configuration are applied.

Options

juniper-enhanced Enables Enhanced Web Filtering (EWF) on the device.

base-filter A base filter is an object that contains a category-action pair for all categories defined in the category file.

block-message	Juniper enhanced block message settings.
category	Select a custom URL category list you created (custom objects) for filtering against.
custom-block-message	Enter a custom message to be sent when HTTP requests are blocked.
default	Specify an action for the profile, for requests that experience internal errors in the Web filtering module.
fallback-settings	Fallback settings helps the system how to handle errors.
no-safe-search	Do not perform safe-search for Juniper enhanced protocol. Safe-search redirect supports HTTP only. Therefore it is not possible to generate a redirect response for HTTPS search URLs. Safe-search redirects can be disabled by using the CLI option `no-safe-search`.
quarantine-custom-message	Juniper enhanced quarantine custom message.
quarantine-message	Juniper enhanced quarantine message settings.
server	Set server parameters by entering the server name or IP address.
site-reputation-action	Specify the action to be taken depending on the site reputation returned for all types of URLs whether it is categorized or uncategorized.
timeout	Enter a timeout limit for requests. Once this limit is reached, fail mode settings are applied. • **Range:** 1 through 1800
juniper-local	Enables Juniper Networks local URL filtering on the device.
block-message	Juniper local block message settings.
websense-redirect	Web filtering websense redirect engine. Websense occasionally releases new EWF categories. EWF classifies websites into categories according to host, URL, or IP address and performs filtering based on the categories.
type	Type of Web filtering solution or URL filtering solution used by the device.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.3R1.

Support for configuration in tenant systems introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Unified Threat Management Overview

utm default-configuration

zone (System Security Profile)

IN THIS SECTION

- Syntax | **890**
- Hierarchy Level | **890**
- Description | **890**
- Options | **890**
- Required Privilege Level | **891**
- Release Information | **891**

Syntax

```
zone {
    maximum maximum;
    reserved reserved;
}
```

Hierarchy Level

```
[edit system security-profile]
```

Description

Specify the zones that user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The primary administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the primary logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the primary administrator can create security profiles and bind them to logical systems.

Options

maximum A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The

maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **Range:** 0 through 4000

reserved A reserved quota that guarantees that the resource amount specified is always available to the logical system.

- **Default:** 0

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.2.

RELATED DOCUMENTATION

Understanding Logical Systems Security Profiles (Primary Administrators Only) | 71

Example: Configuring Logical Systems Security Profiles (Primary Administrators Only) | 77

5

CHAPTER

Operational Commands

clear class-of-service application-traffic-control counter | 896

clear class-of-service application-traffic-control rate-limiters | 898

clear class-of-service application-traffic-control statistics rule | 900

clear security application-firewall rule-set statistics logical-system | 902

clear security dns-cache | 904

clear security firewall-authentication users | 906

clear security firewall-authentication history | 909

clear security idp attack table | 912

clear security idp counters ips | 913

clear security idp counters pdf-decoder | 915

clear security idp counters ssl-inspection | 917

clear security idp counters memory | 919

clear security idp counters memory | 921

clear security idp counters tcp-reassembler | 922

clear security idp counters application-identification | 924

clear security idp counters action | 926

clear security idp counters dfa | 928

clear security idp counters flow | 929

clear security idp counters log | 931

clear security idp counters http-decoder | 933

clear security idp counters packet-log | 934

clear security idp counters packet | 936

clear security idp counters policy-manager | 938

clear security flow session tenant | 940

clear services user-identification logical-domain-identity-management counters | 942

request security datapath-debug capture start | 944

request security datapath-debug capture stop | 945

set chassis cluster cluster-id node node-number reboot | 947

show chassis cluster status | 949

show class-of-service application-traffic-control rate-limiters | 954

show log | 966

show route tenant | 974

show security application-firewall rule-set | 977

show security application-firewall rule-set logical-system | 982

show security application-tracking counters | 987

show security alg status logical-system | 989

show security datapath-debug capture | 994

show security datapath-debug counter | 996

show security dns-cache | 998

show security dynamic-address | 1001

show security firewall-authentication history | 1011

show security firewall-authentication users | 1015

show security flow session | 1019

show security flow session tenant | 1034

show security idp logical system | 1037

show security idp attack table | 1039

show security idp counters action | 1042

show security idp counters application-identification | 1046

show security idp counters memory | 1053

show security idp counters ssl-inspection | 1057

show security idp counters pdf-decoder | 1061

show security idp counters log | 1065

show security idp counters ips | 1073

show security idp counters dfa | 1081

show security idp counters flow | 1084

show security idp counters http-decoder | 1101

show security idp counters packet-log | 1105

show security idp counters packet | 1120

show security idp counters policy-manager | 1130

show security idp counters tcp-reassembler | 1133

show security idp logical-system policy-association | 1141

show security idp policies | 1143

show security idp policy-commit-status | 1145

show security idp policy-templates-list | 1148

show security idp security-package-version | 1150

show security ike security-associations | 1153

show security ipsec security-associations | 1171

show security log report | 1209

show security match-policies | 1211

show security nat destination rule | 1221

show security nat destination summary | 1226

show security nat source rule | 1230

show security nat source summary | 1237

show security nat static rule | 1245

show security policies | 1252

show security screen statistics | 1276

show services user-identification logical-domain-identity-management | 1292

show system security-profile | 1296

show system security-profile secure-wire | 1305

show system security-profile scheduler | 1310

show system security-profile security-log-stream-number detail | 1315

show system security-profile security-log-stream-number | 1318

show system security-profile security-log-stream-number summary | 1322

show security softwires | 1325

show security softwires map-e confidentiality status | 1328

show security softwires map-e domain | 1329

show security zones | 1332

clear class-of-service application-traffic-control counter

IN THIS SECTION

- Syntax | 896
- Description | 896
- Options | 897
- Required Privilege Level | 897
- Output Fields | 897
- Sample Output | 897
- Release Information | 898

Syntax

```
clear class-of-service application-traffic-control counter
<all-logical-systems-tenants>
<logical-system logical-system-name>
<node>
<root-logical-system>
<tenant tenant-name>
```

Description

Clears the application traffic control counters.

Options

all-logical-systems-tenants	(Optional) Clears the application traffic control counters of logical systems and tenant systems.
logical-system *logical-system-name*	(Optional) Clears the application traffic control counters of a specified logical system.
node	(Optional) Clears the application traffic control counters of a specified node.
root-logical-system	(Optional) Clears the application traffic control counters of the root logical-system (default).
tenant *tenant-name*	(Optional) Clears the application traffic control counters of a specified tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

clear class-of-service application-traffic-control counter

```
user@host>clear class-of-service application-traffic-control counter
```

Release Information

Command introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

show class-of-service application-traffic-control counter

clear class-of-service application-traffic-control rate-limiters

IN THIS SECTION

- Syntax | 898
- Description | 898
- Required Privilege Level | 899
- Output Fields | 899
- Sample Output | 899
- Release Information | 899

Syntax

```
clear class-of-service application-traffic-control rate-limiters
```

Description

Clears the application traffic control rate limiters information.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

clear class-of-service application-traffic-control rate-limiters

```
user@host>clear class-of-service application-traffic-control rate-limiters
```

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

show class-of-service application-traffic-control rate-limiters | 954

clear class-of-service application-traffic-control statistics rule

IN THIS SECTION

- Syntax | 900
- Description | 900
- Options | 901
- Required Privilege Level | 901
- Output Fields | 901
- Sample Output | 901
- Release Information | 902

Syntax

```
clear class-of-service application-traffic-control statistics rule
<all-logical-systems-tenants>
<logical-system logical-system-name>
<node>
<root-logical-system>
<tenant tenant-name>
```

Description

Clears the application traffic control rule hit statistics. AppQoS counters identifying rule hits.

Options

all-logical-systems-tenants	Clears the application traffic control rule hit statistics of logical systems and tenant systems.
logical-system *logical-system-name*	(Optional) Clears the application traffic control rule hit statistics of a specified logical-system.
node	(Optional) Clears the application traffic control rule hit statistics of a specified node.
root-logical-system	(Optional) Clears the application traffic control rule hit statistics of the root logical system (default).
tenant *tenant-name*	(Optional) Clears the application traffic control rule hit statistics of a specified tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Sample Output

clear class-of-service application-traffic-control statistics rule

```
user@host> clear class-of-service application-traffic-control statistics rule
```

Release Information

Command introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

show class-of-service application-traffic-control statistics rule

clear security application-firewall rule-set statistics logical-system

IN THIS SECTION

- Syntax | 902
- Description | 903
- Options | 903
- Required Privilege Level | 903
- Output Fields | 904
- Release Information | 904

Syntax

The primary, or root, administrator can issue the following statements:

```
clear security application-firewall rule-set statistics [logical-system logical-system-name |
all | root-logical-system]
```

The user logical system administrator can issue the following statement:

```
clear security application-firewall rule-set statistics all
```

Description

Clear all security application firewall rule set statistics.

NOTE: User logical system administrators can clear statistics only for the logical systems they can access. For information about primary and user administrator roles in logical systems, see Understanding the Primary Logical Systems and the Primary Administrator Role.

Starting in Junos OS Release 18.2R1 application firewall (AppFW) functionality is deprecated. As a part of this change, the `[edit security application-firewall]` hierarchy and all the configuration options under this hierarchy are deprecated— rather than immediately removed—to provide backward compatibility and a chance to bring your configuration into compliance with the new configuration.

Options

logical-system-name—Name of a specific logical system.

`all`—(default) Clear all rule set statistics for a specific logical system or all logical systems.

`root-logical-system`—Clear application firewall rule set statistics on the root logical system (primary administrator only).

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

show security application-firewall rule-set logical-system

clear security dns-cache

IN THIS SECTION

- Syntax | **904**
- Description | **905**
- Options | **905**
- Required Privilege Level | **905**
- Release Information | **905**

Syntax

```
clear security dns-cache <dns-name dns-name>
```

Description

Reset DNS cache information.

> **NOTE**: This command is only available to the primary administrator on devices that are configured for logical systems. This command is not available in user logical systems or on devices that are not configured for logical systems.

Options

- `dns-name`—Clear DNS cache information for the specified name.

Required Privilege Level

clear

Release Information

Command introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

show security dns-cache | 998

Understanding the Primary Logical Systems and the Primary Administrator Role | 23

clear security firewall-authentication users

IN THIS SECTION

- Syntax | 906
- Description | 906
- Options | 907
- Required Privilege Level | 907
- Output Fields | 907
- Sample Output | 908
- Sample Output | 908
- Release Information | 908

Syntax

```
clear security firewall-authentication users
<node (node-id | all | local | primary)>
<address>
<identifier>
<logical-system (logical-system-name | all)>
<root-logical-system (address | auth-type | from-zone | identifier | tenant | to-zone>
<tenant (tenant-name |all)>
```

Description

Clear firewall authentication tables for all users.

Options

- node—(Optional) For chassis cluster configurations, clear firewall authentication details for all users on a specific node.
  - *node-id* —Identification number of the node. It can be 0 or 1.
  - all —Clear all nodes.
  - local —Clear the local node.
  - primary—Clear the primary node.
- address—Display authentication entries based on ip address.
- identifier—Display authentication entries by id.
- logical-system—Display firewall authentication tables based on logical system name.
- node—(Optional) For chassis cluster configurations, display firewall authentication details for all users on a specific node.
  - *node-id*—Identification number of the node. It can be 0 or 1.
  - all—Display information about all nodes.
  - local—Display information about the local node.
  - primary—Display information about the primary node.
- root-logical-system—Display firewall authentication tables for root logical system.
- tenant—Display firewall authentication tables based on tenant name.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

This command produces no output.

Sample Output

clear security firewall-authentication users

```
user@host> clear security firewall-authentication users node 1
node0:
--------------------------------------------------------------------------
node1:
--------------------------------------------------------------------------
```

Sample Output

clear security firewall-authentication users node 1

```
user@host> clear security firewall-authentication users node 1
node1:
--------------------------------------------------------------------------
```

clear security firewall-authentication users tenant all

```
user@host> clear security firewall-authentication users tenant all
```

Release Information

Command introduced in Junos OS Release 8.5. The `node` options added in Junos OS Release 9.0.

The `tenant` option introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Firewall User Authentication Overview

show security firewall-authentication users

clear security firewall-authentication history

IN THIS SECTION

- Syntax | 909
- Description | 909
- Options | 910
- Required Privilege Level | 910
- Output Fields | 910
- Sample Output | 911
- Sample Output | 911
- Release Information | 911

Syntax

```
clear security firewall-authentication history
<node (node-id | all | local | primary)>
<address>
<identifier>
<logical-system (logical-system-name | all)>
<root-logical-system (address | auth-type | from-zone | identifier | tenant | to-zone>
<tenant (tenant-name |all)>
```

Description

Clear all firewall authentication history information.

Options

- node—(Optional) For chassis cluster configurations, clear all firewall authentication history on a specific node (device) in the cluster.
  - *node-id* —Identification number of the node. It can be 0 or 1.
  - all —Clear all nodes.
  - local —Clear the local node.
  - primary—Clear the primary node.
- address—Display authentication entries based on ip address.
- identifier—Display authentication entries by id.
- logical-system—Display firewall authentication tables based on logical system name.
- node—(Optional) For chassis cluster configurations, display firewall authentication details for all users on a specific node.
  - *node-id*—Identification number of the node. It can be 0 or 1.
  - all—Display information about all nodes.
  - local—Display information about the local node.
  - primary—Display information about the primary node.
- root-logical-system—Display firewall authentication tables for root logical system.
- tenant—Display firewall authentication tables based on tenant name.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

This command produces no output.

Sample Output

clear security firewall-authentication history

```
user@host> clear security firewall-authentication history
node0:
--------------------------------------------------------------------------
node1:
--------------------------------------------------------------------------
```

Sample Output

clear security firewall-authentication history node 1

```
user@host> clear security firewall-authentication history node 1
node1:
--------------------------------------------------------------------------
```

clear security firewall-authentication history tenant all

```
user@host> clear security firewall-authentication history tenant all
```

Release Information

Command introduced in Junos OSRelease 8.5. The `node` options added in Junos OS Release 9.0.

The `tenant` option introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

> *Firewall User Authentication Overview*
>
> *show security firewall-authentication history*

clear security idp attack table

IN THIS SECTION

- Syntax | 912
- Description | 912
- Options | 912
- Required Privilege Level | 913
- Output Fields | 913
- Release Information | 913

Syntax

```
clear security idp attack table
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Clears the details of the IDP attack table.

Options

none	Clears the details of the IDP attack table.
logical-system *logical-system-name*	(Optional) Clears the details of the IDP attack table for a specific logical system.
logical-system all	(Optional) Clears the details of the IDP attack table for all logical systems.

tenant *tenant-name* (Optional) Clears the details of the IDP attack table for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show security idp attack table

clear security idp counters ips

IN THIS SECTION

- Syntax | **914**
- Description | **914**
- Options | **914**

- Required Privilege Level | **914**
- Output Fields | **915**
- Release Information | **915**

Syntax

```
clear security idp counters ips
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Resets all the IPS counter values.

Options

none	Resets all the IPS counter values.
logical-system *logical-system-name*	(Optional) Resets all the IPS counter values for a specific logical system.
logical-system all	(Optional) Resets all the IPS counter values for all logical systems.
tenant *tenant-name*	(Optional) Resets all the IPS counter values for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

ips

show security idp counters ips

clear security idp counters pdf-decoder

IN THIS SECTION

- Syntax | **916**
- Description | **916**
- Options | **916**
- Required Privilege Level | **916**
- Output Fields | **916**
- Release Information | **917**

Syntax

```
clear security idp counters pdf-decoder
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Resets all the PDF-Decoder enabled sessions.

Options

none	Resets all the PDF-Decoder enabled sessions.
logical-system *logical-system-name*	(Optional) Resets all the PDF-Decoder enabled sessions for a specific logical system.
logical-system all	(Optional) Resets all the PDF-Decoder enabled sessions for all logical systems.
tenant *tenant-name*	(Optional) Resets all the PDF-Decoder enabled sessions for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced for user logical systems in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show security idp counters pdf-decoder | 1061

clear security idp counters ssl-inspection

IN THIS SECTION

- Syntax | 917
- Description | 918
- Options | 918
- Required Privilege Level | 918
- Output Fields | 918
- Release Information | 918

Syntax

```
clear security idp counters ssl-inspection
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Resets all the session encryption and decryption values for IDP counters.

Options

none	Resets all the session encryption and decryption values for IDP counters.
logical-system *logical-system-name*	(Optional) Resets all the session encryption and decryption values for IDP counters for a specific logical system.
logical-system all	(Optional) Resets all the session encryption and decryption values for IDP counters for all logical systems.
tenant *tenant-name*	(Optional) Resets all the session encryption and decryption values for IDP counters for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced for user logical systems in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show security idp counters ssl-inspection | 1057

clear security idp counters memory

IN THIS SECTION

- Syntax | 919
- Description | 919
- Options | 919
- Required Privilege Level | 920
- Output Fields | 920
- Release Information | 920

Syntax

```
clear security idp counters memory
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Resets all the memory allocation and reallocation counter values.

Options

none	Resets all the memory allocation and reallocation counter values.

logical-system *logical-system-name*	(Optional) Resets all the memory allocation and reallocation counter values for a specific logical system.
logical-system all	(Optional) Resets all the memory allocation and reallocation counter values for all logical systems.
tenant *tenant-name*	(Optional) Resets all the memory allocation and reallocation counter values for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced for user logical systems in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show security idp counters memory | 1053

clear security idp counters memory

IN THIS SECTION

- Syntax | 921
- Description | 921
- Options | 921
- Required Privilege Level | 922
- Output Fields | 922
- Release Information | 922

Syntax

```
clear security idp counters memory
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Resets all the memory allocation and reallocation counter values.

Options

none	Resets all the memory allocation and reallocation counter values.
logical-system *logical-system-name*	(Optional) Resets all the memory allocation and reallocation counter values for a specific logical system.

logical-system all	(Optional) Resets all the memory allocation and reallocation counter values for all logical systems.
tenant *tenant-name*	(Optional) Resets all the memory allocation and reallocation counter values for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced for user logical systems in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show security idp counters memory | 1053

clear security idp counters tcp-reassembler

IN THIS SECTION

- Syntax | 923
- Description | 923

- Options | 923
- Required Privilege Level | 923
- Output Fields | 924
- Release Information | 924

Syntax

```
clear security idp counters tcp-reassembler
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Resets all the TCP reassembler counter values.

Options

none	Resets all the TCP reassembler counter values.
logical-system *logical-system-name*	(Optional) Resets all the TCP reassembler counter values for a specific logical system.
logical-system all	(Optional) Resets all the TCP reassembler counter values for all logical systems.
tenant *tenant-name*	(Optional) Resets all the TCP reassembler counter values for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

re-assembler

show security idp counters tcp-reassembler

clear security idp counters application-identification

IN THIS SECTION

- Syntax | **925**
- Description | **925**
- Options | **925**
- Required Privilege Level | **925**
- Output Fields | **925**
- Release Information | **926**

Syntax

```
clear security idp counters application-identification
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Resets all the application identification counter values.

Options

none	Resets all the application identification counter values.
logical-system *logical-system-name*	(Optional) Resets all the application identification counter values for a specific logical system.
logical-system all	(Optional) Resets all the application identification counter values for all logical systems.
tenant *tenant-name*	(Optional) Resets all the application identification counter values for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

application-identification

show security idp counters application-identification

clear security idp counters action

IN THIS SECTION

- Syntax | **926**
- Description | **927**
- Options | **927**
- Required Privilege Level | **927**
- Output Fields | **927**
- Release Information | **927**

Syntax

```
clear security idp counters action
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Resets all the action counter values.

Options

none	Resets all the action counter values.
logical-system *logical-system-name*	(Optional) Resets all the action counter values for a specific logical system.
logical-system all	(Optional) Resets all the action counter values for all logical systems.
tenant *tenant-name*	(Optional) Resets all the action counter values for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced for user logical systems in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show security idp counters action | 1042

clear security idp counters dfa

IN THIS SECTION

- Syntax | 928
- Description | 928
- Options | 928
- Required Privilege Level | 929
- Output Fields | 929
- Release Information | 929

Syntax

```
clear security idp counters dfa
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Resets all the DFA counter values.

Options

none	Resets all the DFA counter values.
logical-system *logical-system-name*	(Optional) Resets all the DFA counter values for a specific logical system.
logical-system all	(Optional) Resets all the DFA counter values for all logical systems.

tenant *tenant-name*	(Optional) Resets all the DFA counter values for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show security idp counters dfa

clear security idp counters flow

IN THIS SECTION

- Syntax | **930**
- Description | **930**
- Required Privilege Level | **930**

Output Fields | 930
Release Information | 930

Syntax

```
clear security idp counters flow
clear security idp counters flow logical-system logical-system
```

Description

Reset all the IDP flow-related counter values.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 9.2.

Command introduced for user logical systems in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

flow (Security IDP)

show security idp counters flow

clear security idp counters log

IN THIS SECTION

- Syntax | **931**
- Description | **931**
- Options | **932**
- Required Privilege Level | **932**
- Output Fields | **932**
- Release Information | **932**

Syntax

```
clear security idp counters log
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Resets all the IDP log counter values.

Options

none	Resets all the IDP log counter values.
logical-system *logical-system-name*	(Optional) Resets all the IDP log counter values for a specific logical system.
logical-system all	(Optional) Resets all the IDP log counter values for all logical systems.
tenant *tenant-name*	(Optional) Resets all the IDP log counter values for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

event-rate

show security idp counters log

clear security idp counters http-decoder

IN THIS SECTION

- Syntax | 933
- Description | 933
- Options | 933
- Required Privilege Level | 934
- Output Fields | 934
- Release Information | 934

Syntax

```
clear security idp counters http-decoder
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Resets all the HTTP decoder counter values.

Options

none	Resets all the HTTP decoder counter values.
logical-system *logical-system-name*	(Optional) Resets all the HTTP decoder counter values for a specific logical system.
logical-system all	(Optional) Resets all the HTTP decoder counter values for all logical systems.

tenant *tenant-name* (Optional) Resets all the HTTP decoder counter values for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 11.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show security idp counters http-decoder

clear security idp counters packet-log

IN THIS SECTION

- Syntax | 935
- Description | 935

- Options | 935
- Required Privilege Level | 935
- Output Fields | 936
- Release Information | 936

Syntax

```
clear security idp counters packet-log
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Resets all the IDP counters value for packet log.

Options

none	Resets all the IDP counters value for packet log.
logical-system *logical-system-name*	(Optional) Resets all the IDP counters value for packet log for a specific logical system.
logical-system all	(Optional) Resets all the IDP counters value for packet log for all logical systems.
tenant *tenant-name*	(Optional) Resets all the IDP counters value for packet log for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced for user logical systems in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show security idp counters packet-log | **1105**

clear security idp counters packet

IN THIS SECTION

- Syntax | **937**
- Description | **937**
- Options | **937**
- Required Privilege Level | **937**
- Output Fields | **937**
- Release Information | **938**

Syntax

```
clear security idp counters packet
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Resets all the IDP packet counter values.

Options

none	Resets all the IDP packet counter values.
logical-system *logical-system-name*	(Optional) Resets all the IDP packet counter values for a specific logical system.
logical-system all	(Optional) Resets all the IDP packet counter values for all logical systems.
tenant *tenant-name*	(Optional) Resets all the IDP packet counter values for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show security idp counters packet

clear security idp counters policy-manager

IN THIS SECTION

- Syntax | **938**
- Description | **939**
- Options | **939**
- Required Privilege Level | **939**
- Output Fields | **939**
- Release Information | **939**

Syntax

```
clear security idp counters policy-manager
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Resets all the IDP policies counter values.

Options

none	Resets all the IDP policies counter values.
logical-system *logical-system-name*	(Optional) Resets all the IDP policies counter values for a specific logical system.
logical-system all	(Optional) Resets all the IDP policies counter values for all logical systems.
tenant *tenant-name*	(Optional) Resets all the IDP policies counter values for a specific tenant system.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show security idp counters policy-manager

clear security flow session tenant

IN THIS SECTION

- Syntax | **940**
- Description | **940**
- Options | **941**
- Required Privilege Level | **941**
- Output Fields | **941**
- Sample Output | **941**
- Release Information | **942**

Syntax

```
clear security flow session tenant (tenant-name| all)
```

Description

Clears the information about the currently active security flow sessions of the tenant systems on the device. You can either clear the currently active security flow sessions for a specific tenant system or for all the tenant systems.

Options

tenant-name	Name of the tenant system.
all	Clears the security flow session information for all the tenant systems.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security flow session tenant T1

```
user@host> clear security flow session tenant T1
0 active sessions cleared
1 active sessions cleared
0 active sessions cleared
0 active sessions cleared
```

clear security flow session tenant all

```
user@host> clear security flow session tenant all
0 active sessions cleared
2 active sessions cleared
0 active sessions cleared
0 active sessions cleared
```

Release Information

Command introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

show security flow session tenant | 1034

clear services user-identification logical-domain-identity-management counters

IN THIS SECTION

- Syntax | 942
- Description | 942
- Options | 943
- Required Privilege Level | 943
- Sample Output | 943
- Release Information | 943

Syntax

```
clear services user-identification logical-domain-identity-management counters
```

Description

Clears logical domain identity management information.

Options

counters Clears logical domain identity management query counters.

Required Privilege Level

clear

Sample Output

clear services user-identification logical-domain-identity-management counters

```
user@host> clear services user-identification logical-domain-identity-management counters
node0:
--------------------------------------------------------------------------
warning: "There is no logical-domain-identity-management."
```

Release Information

Command introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

show services user-identification logical-domain-identity-management | 1292

request security datapath-debug capture start

IN THIS SECTION

- Syntax | 944
- Description | 944
- Required Privilege Level | 944
- Output Fields | 945
- Sample Output | 945
- Release Information | 945

Syntax

```
request security datapath-debug capture start
```

Description

Start the data path debugging capture.

NOTE: Data path debugging is supported on SRX1400, SRX3400, SRX3600, SRX5400, SRX5600, and SRX5800.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security datapath-debug capture start

```
user@host> request  security datapath-debug capture start
datapath-debug capture started on file
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Understanding Data Path Debugging for Logical Systems | 520

request security datapath-debug capture stop

IN THIS SECTION

- Syntax | 946
- Description | 946
- Required Privilege Level | 946
- Output Fields | 946
- Sample Output | 946
- Release Information | 946

Syntax

```
request security datapath-debug capture stop
```

Description

Stop the data path debugging capture.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security datapath-debug capture stop

```
user@host> request  security datapath-debug capture stop
datapath-debug capture successfully stopped, use show security datapath-debug
capture to view
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

Understanding Data Path Debugging for Logical Systems | 520

set chassis cluster cluster-id node node-number reboot

IN THIS SECTION

- Syntax | 947
- Description | 947
- Options | 948
- Required Privilege Level | 948
- Output Fields | 948
- Release Information | 948

Syntax

```
set chassis cluster cluster-id cluster-id node  node-number reboot
```

Description

Sets the chassis cluster identifier (ID) and node ID on each device, and reboots the devices to enable clustering. 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. A cluster ID greater than 15 can only be set when the fabric and control link interfaces are connected back-to-back.

NOTE: If you have a cluster set up and running with an earlier release of Junos OS, you can upgrade to Junos OS Release 12.1X45-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.

Options

cluster-id *cluster-id*	Identifies the cluster within the Layer 2 domain. • **Range:** 0 through 255
node *node*	Identifies a node within a cluster. • **Range:** 0 through 1
reboot	reboot the specified devices to configure the chassis cluster.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Release Information

Support for extended cluster identifiers (more than 15 identifiers) added in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

- Example: Setting the Chassis Cluster Node ID and Cluster ID
- Understanding the Interconnect Logical System and Logical Tunnel Interfaces
- Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (Primary Administrators Only)
- Disabling a Chassis Cluster
- set chassis cluster disable reboot

show chassis cluster status

IN THIS SECTION

- Syntax | 949
- Description | 950
- Options | 950
- Required Privilege Level | 950
- Output Fields | 950
- Sample Output | 952
- Sample Output | 952
- Sample Output | 953
- Release Information | 953

Syntax

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

Description

Display the current status of the Chassis Cluster. You can use this command to check the status of chassis cluster nodes, redundancy groups, and failover status.

Options

- none—Display the status of all redundancy groups in the chassis cluster.
- `redundancy-group` *group-number*—(Optional) Display the status of the specified redundancy group.

Required Privilege Level

view

Output Fields

Table 41 on page 950 lists the output fields for the `show chassis cluster status` command. Output fields are listed in the approximate order in which they appear.

Table 41: show chassis cluster status Output Fields

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

Table 41: show chassis cluster status Output Fields *(Continued)*

Field Name	Field Description
`Status`	State of the redundancy group (`Primary`, `Secondary`, `Lost`, or `Unavailable`). • `Primary`—Redundancy group is active and passing traffic. • `Secondary`—Redundancy group is passive and not passing traffic. • `Lost`—Node loses contact with the other node through the control link. Most likely to occur when both nodes are in a cluster and there is a control link failure, one node cannot exchange heartbeats, or when the other node is rebooted. • `Unavailable`—Node has not received a single heartbeat over the control link from the other node since the other node booted up. Most likely to occur when one node boots up before the other node, or if only one node is present in the cluster.
`Preempt`	• `Yes`: Primary state can be preempted based on priority. • `No`: Change in priority will not preempt the primary state.
`Manual failover`	• `Yes`: Primary state is set manually through the CLI with the `request chassis cluster failover node` or `request chassis cluster failover redundancy-group` command. This overrides `Priority` and `Preempt`. • `No`: Primary state is not set manually through the CLI.
Monitor-failures	• `None`: Cluster working properly. • `Monitor Failure code`: Cluster is not working properly and the respective failure code is displayed.

Sample Output

show chassis cluster status

```
user@host> show chassis cluster status

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

Cluster ID: 1
Node   Priority Status         Preempt Manual   Monitor-failures

Redundancy group: 0 , Failover count: 1
node0  200      primary        no      no       None
node1  1        secondary      no      no       None

Redundancy group: 1 , Failover count: 1
node0  101      primary        no      no       None
node1  1        secondary      no      no       None
```

Sample Output

show chassis cluster status with preemptive delay

```
user@host> show chassis cluster status

Cluster ID: 1
Node   Priority Status         Preempt Manual   Monitor-failures
Redundancy group: 0, Failover count: 1
node0  200      primary        no      no       None
node1  100      secondary      no      no       None
```

```
Redundancy group: 1, Failover count: 3
node0  200      primary-preempt-hold       yes     no       None node1  100
secondary                    yes     no       None
```

Sample Output

show chassis cluster status redundancy-group 1

```
user@host> show chassis cluster status redundancy-group 1

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

Cluster ID: 1
Node   Priority Status         Preempt Manual   Monitor-failures

Redundancy group: 1 , Failover count: 1
node0  101      primary        no      no       None
node1  1        secondary      no      no       None
```

Release Information

Support for monitoring failures added in Junos OS Release 12.1X47-D10.

RELATED DOCUMENTATION

- redundancy-group (Chassis Cluster)
- clear chassis cluster failover-count
- request chassis cluster failover node

request chassis cluster failover reset

show class-of-service application-traffic-control rate-limiters

IN THIS SECTION

- Syntax | 954
- Description | 954
- Options | 955
- Additional Information | 955
- Required Privilege Level | 955
- Output Fields | 956
- Sample Output | 958
- Release Information | 966

Syntax

```
show class-of-service application-traffic-control rate-limiters (rl-all | rl-name | summary)
<all-logical-systems-tenants>
<logical-system (logical-system-name | all)>
<root-logical-system>
<tenant (tenant-name | all)>
```

Description

Displays the application traffic control rate limiters information.

Options

rl-all	(Optional) Displays all application traffic control rate limiters information.
rl-name	(Optional) Displays the application traffic control rate limiters information by name.
summary	(Optional) Displays the application traffic control rate limiters information summary.
all-logical-systems-tenants	(Optional) Displays application traffic control rate limiters information of all logical systems and tenant systems.
logical-system *logical-system-name*	(Optional) Displays application traffic control rate limiters information of a specified logical system.
logical-system all	(Optional) Displays application traffic control rate limiters information of all logical systems.
root-logical-system	(Optional) Displays application traffic control rate limiters information of the root logical system as default.
tenant *tenant-name*	(Optional) Displays application traffic control rate limiters information of a specified tenant system.
logical-system all	(Optional) Displays application traffic control rate limiters information of all tenant systems.

Additional Information

Required Privilege Level

view

Output Fields

Table 42 on page 956 lists the output fields for the `show class-of-service application-traffic-control rate-limiters` command. Output fields are listed in the approximate order in which they appear.

Table 42: show class-of-service application-traffic-control rate-limiters Output Fields

Field Name	Field Description
`LSYS/Tenant`	Name of the logical system or tenant system.
`ID`	ID of the rate limiter.
`BW-Limit`	The bandwidth of the rate limiter.
`Burst-Limit`	The burst limit of the rate limiter.
`PerIP-Burst-Limit`	Burst limit per IP.
`Strict-BW-hit`	Number of strict bandwidth hit.
`Total Session processed`	Total number of sessions processed.
`Total PerIP_RLs used`	Total number of rate limiters used per IP.
`PerIP-RLs`	Active rate limiters per IP.
`Sessions`	Number of active sessions.
`Pkts-Trans`	Number of packets transferred.
`Pkts-Drop`	Number of packets dropped.
`Rate-Trans(bps)`	Rate at which packets are transferred.

Table 42: show class-of-service application-traffic-control rate-limiters Output Fields *(Continued)*

Field Name	Field Description
Rate-Drop(bps)	Rate at which packets are dropped.
Rate-LP(bps)	Rate at which packets with LP are dropped.
RL-Name	Name of the rate limiter
Counter type	Type of the counter.
Rate Limiters	Number of rate limiters.
Sessions Processed	Number of sessions processed.
PerIP RLimiters used	Number of rate limiters used per IP.
Active Sessions	Number of active sessions.
Active perIP RLimiters	Number of active rate limiters per IP.
Sessions discarded strict-bw	Number of sessions discarded based on strict bandwidth.
Sessions discarded no perIP	Number of sessions discarded per IP.
Sessions discarded no memory	Number of sessions discarded due to no memory.
Maximum Rate-Limiters	Number of maximum rate limiters.
Maximum perIP-RLimiters	Number of maximum rate limiters per IP.
Total Pkts Tx	Number of total packets transferred.

Table 42: show class-of-service application-traffic-control rate-limiters Output Fields *(Continued)*

Field Name	Field Description
Total Pkts discarded	Number of total packets discarded.
Total Pkts high LP	Number of total high LP packets.
Rate Tx (bps)	Rate at which packets are transferred.
Rate Discard (bps)	Rate at which packets are dropped.
Rate LP (bps)	Rate at which packets with LP are dropped.

Sample Output

show class-of-service application-traffic-control rate-limiters rl-all

```
root@host> show class-of-service application-traffic-control rate-limiters rl-all
LSYS/Tenant ID   BW-Limit   PerIP-RLs Sessions Pkts-Trans Pkts-Drop Pkts-LP Rate-Trans(bps) Rate-
Drop(bps) Rate-LP(bps) RL-Name
root        3    1000       0         0        0          0         0       0
0              0           R1
```

show class-of-service application-traffic-control rate-limiters rl-name R1

```
root@host> show class-of-service application-traffic-control rate-limiters rl-name R1
Logical System: root-logical-system
Counter type                            Value
Name                                    R1
BW-Mode                                 -
ID                                      3
BW-Limit                                1000
Burst-Limit                             128000
```

```
PerIP-BW-Limit                          0
PerIP-Burst-Limit                       0
Strict-BW-hit                           0
Total Session processed                 0
Total PerIP_RLs used                    0
Active-Sessions                         0
Active-PerIP-RL                         0
Total-Pkt-transmit                      0
Total-Pkt-Drop                          0
Total-Pkt-LP                            0
Total-Rate-Transmit(bps)                0
Total-Rate-Drop(bps)                    0
Total-Rate-LP(bps)                      0
RL-Pkt-transmit                         0
RL-Pkt-Drop                             0
RL-Pkt-LP                               0
RL-Rate-Transmit(bps)                   0
RL-Rate-Drop(bps)                       0
RL-Rate-LP(bps)                         0
```

show class-of-service application-traffic-control rate-limiters rl-all all-logical-systems-tenants

```
root@host> show class-of-service application-traffic-control rate-limiters rl-all all-logical-
systems-tenants
LSYS/Tenant ID BW-Limit PerIP-RLs Sessions Pkts-Trans Pkts-Drop Pkts-LP  Rate-Trans(bps) Rate-
Drop(bps)  Rate-LP(bps) RL-Name
root        3  1000     0         0        0          0         0        0
0               0           R1
TSYS1       1  1000     0         0        0          0         0        0
0               0           rt11
TSYS2       2  1500     0         0        0          0         0        0
0               0           rt12
```

show class-of-service application-traffic-control rate-limiters rl-all logical-system all

```
root@host> show class-of-service application-traffic-control rate-limiters rl-all logical-system
all
LSYS/Tenant ID  BW-Limit PerIP-RLs Sessions Pkts-Trans Pkts-Drop Pkts-LP Rate-Trans(bps) Rate-
Drop(bps) Rate-LP(bps) RL-Name
```

```
root        3   1000     0         0        0          0         0       0
0              0           R1
```

show class-of-service application-traffic-control rate-limiters rl-all tenant all

```
root@host> show class-of-service application-traffic-control rate-limiters rl-all tenant all
LSYS/Tenant ID BW-Limit  PerIP-RLs Sessions Pkts-Trans Pkts-Drop Pkts-LP Rate-Trans(bps) Rate-
Drop(bps)  Rate-LP(bps) RL-Name
TSYS1         1  1000      0         0        0          0         0       0
0               0            rt11
TSYS2         2  1500      0         0        0          0         0       0
0               0            rt12
```

show class-of-service application-traffic-control rate-limiters rl-all tenant TSYS1

```
root@host> show class-of-service application-traffic-control rate-limiters rl-all tenant TSYS1
LSYS/Tenant ID BW-Limit PerIP-RLs Sessions Pkts-Trans   Pkts-Drop  Pkts-LP  Rate-Trans(bps) Rate-
Drop(bps) Rate-LP(bps) RL-Name
TSYS1       1  1000     0         0        0            0          0        0
0              0           rt11
```

show class-of-service application-traffic-control rate-limiters summary all-logical-systems-tenants

```
root@host> show class-of-service application-traffic-control rate-limiters summary all-logical-
systems-tenants
Logical System: root-logical-system
Counter type                            Value
Rate Limiters                           1
Sessions Processed                      0
PerIP RLimiters used                    0
Active Sessions                         0
Active perIP RLimiters                  0
Sessions discarded strict-bw            0
Sessions discarded no perIP             0
Sessions discarded no memory            0
Maximum Rate-Limiters                   1000
Maximum perIP-RLimiters                 24000
```

```
Total Pkts Tx                          0
Total Pkts discarded                   0
Total Pkts high LP                     0
Rate Tx (bps)                          0
Rate Discard (bps)                     0
Rate LP (bps)                          0

Logical System: TSYS1
Counter type                           Value
Rate Limiters                          1
Sessions Processed                     0
PerIP RLimiters used                   0
Active Sessions                        0
Active perIP RLimiters                 0
Sessions discarded strict-bw           0
Sessions discarded no perIP            0
Sessions discarded no memory           0
Maximum Rate-Limiters                  1000
Maximum perIP-RLimiters                24000
Total Pkts Tx                          0
Total Pkts discarded                   0
Total Pkts high LP                     0
Rate Tx (bps)                          0
Rate Discard (bps)                     0
Rate LP (bps)                          0

Logical System: TSYS2
Counter type                           Value
Rate Limiters                          1
Sessions Processed                     0
PerIP RLimiters used                   0
Active Sessions                        0
Active perIP RLimiters                 0
Sessions discarded strict-bw           0
Sessions discarded no perIP            0
Sessions discarded no memory           0
Maximum Rate-Limiters                  1000
Maximum perIP-RLimiters                24000
Total Pkts Tx                          0
Total Pkts discarded                   0
Total Pkts high LP                     0
Rate Tx (bps)                          0
```

```
Rate Discard (bps)                      0
Rate LP (bps)                           0
```

show class-of-service application-traffic-control rate-limiters summary logical-system all

```
root@host> show class-of-service application-traffic-control rate-limiters summary logical-
system all
Logical System: root-logical-system
Counter type                            Value
Rate Limiters                           1
Sessions Processed                      0
PerIP RLimiters used                    0
Active Sessions                         0
Active perIP RLimiters                  0
Sessions discarded strict-bw            0
Sessions discarded no perIP             0
Sessions discarded no memory            0
Maximum Rate-Limiters                   1000
Maximum perIP-RLimiters                 24000
Total Pkts Tx                           0
Total Pkts discarded                    0
Total Pkts high LP                      0
Rate Tx (bps)                           0
Rate Discard (bps)                      0
Rate LP (bps)                           0
```

show class-of-service application-traffic-control rate-limiters summary tenant TSYS1

```
root@host> show class-of-service application-traffic-control rate-limiters summary tenant TSYS1
Logical System: TSYS1
Counter type                            Value
Rate Limiters                           1
Sessions Processed                      0
PerIP RLimiters used                    0
Active Sessions                         0
Active perIP RLimiters                  0
Sessions discarded strict-bw            0
Sessions discarded no perIP             0
Sessions discarded no memory            0
Maximum Rate-Limiters                   1000
```

```
Maximum perIP-RLimiters                  24000
Total Pkts Tx                            0
Total Pkts discarded                     0
Total Pkts high LP                       0
Rate Tx (bps)                            0
Rate Discard (bps)                       0
Rate LP (bps)                            0
```

show class-of-service application-traffic-control rate-limiters rl-all

```
root@host:TSYS1> show class-of-service application-traffic-control rate-limiters rl-all
LSYS/Tenant ID BW-Limit PerIP-RLs Sessions Pkts-Trans Pkts-Drop Pkts-LP Rate-Trans(bps) Rate-
Drop(bps) Rate-LP(bps) RL-Name
TSYS1         1  1000     0         0        0          0         0       0
0               0           R1
```

show class-of-service application-traffic-control rate-limiters rl-name rt11

```
root@host:TSYS1> show class-of-service application-traffic-control rate-limiters rl-name rt11
Logical System: TSYS1
Counter type                             Value
Name                                     rt11
BW-Mode                                  -
ID                                       1
BW-Limit                                 1000
Burst-Limit                              20000
PerIP-BW-Limit                           0
PerIP-Burst-Limit                        0
Strict-BW-hit                            0
Total Session processed                  0
Total PerIP_RLs used                     0
Active-Sessions                          0
Active-PerIP-RL                          0
Total-Pkt-transmit                       0
Total-Pkt-Drop                           0
Total-Pkt-LP                             0
Total-Rate-Transmit(bps)                 0
Total-Rate-Drop(bps)                     0
Total-Rate-LP(bps)                       0
RL-Pkt-transmit                          0
```

```
RL-Pkt-Drop                              0
RL-Pkt-LP                                0
RL-Rate-Transmit(bps)                    0
RL-Rate-Drop(bps)                        0
RL-Rate-LP(bps)                          0
```

show class-of-service application-traffic-control rate-limiters summary

```
root@host:TSYS1> show class-of-service application-traffic-control rate-limiters summary
Logical System: TSYS1
Counter type                             Value
Rate Limiters                            1
Sessions Processed                       0
PerIP RLimiters used                     0
Active Sessions                          0
Active perIP RLimiters                   0
Sessions discarded strict-bw             0
Sessions discarded no perIP              0
Sessions discarded no memory             0
Maximum Rate-Limiters                    1000
Maximum perIP-RLimiters                  24000
Total Pkts Tx                            0
Total Pkts discarded                     0
Total Pkts high LP                       0
Rate Tx (bps)                            0
Rate Discard (bps)                       0
Rate LP (bps)                            0
```

show class-of-service application-traffic-control rate-limiters rl-all

```
root@host:LSYS1> show class-of-service application-traffic-control rate-limiters rl-all
LSYS/Tenant ID BW-Limit PerIP-RLs Sessions Pkts-Trans Pkts-Drop Pkts-LP Rate-Trans(bps) Rate-
Drop(bps) Rate-LP(bps) RL-Name
LSYS1       1  1000     0         0        0          0         0       0
0              0           rt1
```

show class-of-service application-traffic-control rate-limiters rl-name rt1

```
root@host:LSYS1> show class-of-service application-traffic-control rate-limiters rl-name rt1
Logical System: LSYS1
Counter type                            Value
Name                                    rt1
BW-Mode                                 -
ID                                      1
BW-Limit                                1000
Burst-Limit                             128000
PerIP-BW-Limit                          0
PerIP-Burst-Limit                       0
Strict-BW-hit                           0
Total Session processed                 0
Total PerIP_RLs used                    0
Active-Sessions                         0
Active-PerIP-RL                         0
Total-Pkt-transmit                      0
Total-Pkt-Drop                          0
Total-Pkt-LP                            0
Total-Rate-Transmit(bps)                0
Total-Rate-Drop(bps)                    0
Total-Rate-LP(bps)                      0
RL-Pkt-transmit                         0
RL-Pkt-Drop                             0
RL-Pkt-LP                               0
RL-Rate-Transmit(bps)                   0
RL-Rate-Drop(bps)                       0
RL-Rate-LP(bps)                         0
```

show class-of-service application-traffic-control rate-limiters summary

```
root@host:LSYS1> show class-of-service application-traffic-control rate-limiters summary
Logical System: LSYS1
Counter type                            Value
Rate Limiters                           1
Sessions Processed                      0
PerIP RLimiters used                    0
Active Sessions                         0
Active perIP RLimiters                  0
```

```
Sessions discarded strict-bw              0
Sessions discarded no perIP               0
Sessions discarded no memory              0
Maximum Rate-Limiters                     1000
Maximum perIP-RLimiters                   24000
Total Pkts Tx                             0
Total Pkts discarded                      0
Total Pkts high LP                        0
Rate Tx (bps)                             0
Rate Discard (bps)                        0
Rate LP (bps)                             0
```

Release Information

Command introduced in Junos OS Release 11.4.

`all-logical-systems-tenants`, `logical-system`, `root-logical-system`, and `tenant` options introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

clear class-of-service application-traffic-control rate-limiters | 898

show log

IN THIS SECTION

- Syntax | 967
- Syntax (QFX Series and OCX Series) | 967
- Syntax (TX Matrix Router) | 967
- Description | 967
- Options | 968
- Required Privilege Level | 969

Sample Output | 969
Release Information | 973

Syntax

```
show log
<filename / user <username>>
```

Syntax (QFX Series and OCX Series)

```
show log filename
<device-type (device-id | device-alias)>
```

Syntax (TX Matrix Router)

```
show log
<all-lcc | lcc number | scc>
<filename / user <username>>
```

Description

List log files, display log file contents, or display information about users who have logged in to the router or switch.

NOTE: On MX Series routers, modifying a configuration to replace a service interface with another service interface is treated as a catastrophic event. When you modify a configuration,

the entire configuration associated with the service interface—including NAT pools, rules, and service sets—is deleted and then re-created for the newly specified service interface. If there are active sessions associated with the service interface that is being replaced, these sessions are deleted and the NAT pools are then released, which leads to the generation of the NAT_POOL_RELEASE system log messages. However, because NAT pools are already deleted as a result of the catastrophic configuration change and no longer exist, the NAT_POOL_RELEASE system log messages are not generated for the changed configuration.

Options

none — List all log files.

<all-lcc | lcc *number* | scc> — (Routing matrix only)(Optional) Display logging information about all T640 routers (or line-card chassis) or a specific T640 router (replace *number* with a value from 0 through 3) connected to a TX Matrix router. Or, display logging information about the TX Matrix router (or switch-card chassis).

device-type — (QFabric system only) (Optional) Display log messages for only one of the following device types:

- `director-device`—Display logs for Director devices.
- `infrastructure-device`—Display logs for the logical components of the QFabric system infrastructure, including the diagnostic Routing Engine, fabric control Routing Engine, fabric manager Routing Engine, and the default network Node group and its backup (NW-NG-0 and NW-NG-0-backup).
- `interconnect-device`—Display logs for Interconnect devices.
- `node-device`—Display logs for Node devices.

> **NOTE**: If you specify the `device-type` optional parameter, you must also specify either the `device-id` or `device-alias` optional parameter.

(*device-id* | *device-alias*) — If a device type is specified, display logs for a device of that type. Specify either the device ID or the device alias (if configured).

filename — (Optional) Display the log messages in the specified log file. For the routing matrix, the filename must include the chassis information.

> **NOTE**: The *filename* parameter is mandatory for the QFabric system. If you did not configure a syslog filename, specify the default filename of **messages**.

user <*username*> (Optional) Display logging information about users who have recently logged in to the router or switch. If you include *username*, display logging information about the specified user.

Required Privilege Level

trace

Sample Output

show log

```
user@host> show log
total 57518
-rw-r--r--  1 root  bin      211663 Oct  1 19:44 dcd
-rw-r--r--  1 root  bin      999947 Oct  1 19:41 dcd.0
-rw-r--r--  1 root  bin      999994 Oct  1 17:48 dcd.1
-rw-r--r--  1 root  bin      238815 Oct  1 19:44 rpd
-rw-r--r--  1 root  bin     1049098 Oct  1 18:00 rpd.0
-rw-r--r--  1 root  bin     1061095 Oct  1 12:13 rpd.1
-rw-r--r--  1 root  bin     1052026 Oct  1 06:08 rpd.2
-rw-r--r--  1 root  bin     1056309 Sep 30 18:21 rpd.3
-rw-r--r--  1 root  bin     1056371 Sep 30 14:36 rpd.4
-rw-r--r--  1 root  bin     1056301 Sep 30 10:50 rpd.5
-rw-r--r--  1 root  bin     1056350 Sep 30 07:04 rpd.6
-rw-r--r--  1 root  bin     1048876 Sep 30 03:21 rpd.7
-rw-rw-r--  1 root  bin       19656 Oct  1 19:37 wtmp
```

show log filename

```
user@host> show log rpd
Oct  1 18:00:18 trace_on: Tracing to ?/var/log/rpd? started
Oct  1 18:00:18 EVENT <MTU> ds-5/2/0.0 index 24 <Broadcast PointToPoint Multicast
Oct  1 18:00:18
Oct  1 18:00:19 KRT recv len 56 V9 seq 148 op add Type route/if af 2 addr 192.0.2.21 nhop type
local nhop 192.0.2.21
Oct  1 18:00:19 KRT recv len 56 V9 seq 149 op add Type route/if af 2 addr 192.0.2.22 nhop type
unicast nhop 192.0.2.22
Oct  1 18:00:19 KRT recv len 48 V9 seq 150 op add Type ifaddr index 24 devindex 43
Oct  1 18:00:19 KRT recv len 144 V9 seq 151 op chnge Type ifdev devindex 44
Oct  1 18:00:19 KRT recv len 144 V9 seq 152 op chnge Type ifdev devindex 45
Oct  1 18:00:19 KRT recv len 144 V9 seq 153 op chnge Type ifdev devindex 46
Oct  1 18:00:19 KRT recv len 1272 V9 seq 154 op chnge Type ifdev devindex 47
...

user@host:LSYS1> show log flow_lsys1.log
Nov  7 07:34:09 07:34:09.491800:CID-0:THREAD_ID-00:LSYS_ID-01:RT:got route table lock

Nov  7 07:34:09 07:34:09.491809:CID-0:THREAD_ID-00:LSYS_ID-01:RT:released route table lock

Nov  7 07:34:09 07:34:09.491840:CID-0:THREAD_ID-00:LSYS_ID-01:RT:got route table lock

Nov  7 07:34:09 07:34:09.491841:CID-0:THREAD_ID-00:LSYS_ID-01:RT:released route table lock

Nov  7 07:34:09 07:34:09.491854:CID-0:THREAD_ID-00:LSYS_ID-01:RT:cache final sw_nh 0x0

Nov  7 07:34:09 07:34:09.491868:CID-0:THREAD_ID-00:LSYS_ID-01:RT:got route table lock

Nov  7 07:34:09 07:34:09.491869:CID-0:THREAD_ID-00:LSYS_ID-01:RT:released route table lock

Nov  7 07:34:09 07:34:09.491881:CID-0:THREAD_ID-00:LSYS_ID-01:RT:cache final sw_nh 0x0
user@host:TSYS1> show log flow_tsys1.log
Nov  7 13:21:47 13:21:47.217744:CID-0:THREAD_ID-05:LSYS_ID-32:RT:<192.0.2.0/0-
>198.51.100.0/9011;1,0x0> :

Nov  7 13:21:47 13:21:47.217747:CID-0:THREAD_ID-05:LSYS_ID-32:RT:packet [84] ipid = 39281,
@0x7f490ae56d52

Nov  7 13:21:47 13:21:47.217749:CID-0:THREAD_ID-05:LSYS_ID-32:RT:---- flow_process_pkt: (thd 5):
flow_ctxt type 0, common flag 0x0, mbuf 0x4882b600, rtbl7
```

```
Nov  7 13:21:47 13:21:47.217752:CID-0:THREAD_ID-05:LSYS_ID-32:RT: flow process pak fast ifl 88
in_ifp lt-0/0/0.101

Nov  7 13:21:47 13:21:47.217753:CID-0:THREAD_ID-05:LSYS_ID-32:RT:  lt-0/0/0.101:192.0.2.0-
>198.51.100.0, icmp, (0/0)

Nov  7 13:21:47 13:21:47.217756:CID-0:THREAD_ID-05:LSYS_ID-32:RT: find flow: table 0x11d0a2680,
hash 20069(0xffff), sa 192.0.2.0, da 198.51.100.0, sp 0, d0

Nov  7 13:21:47 13:21:47.217760:CID-0:THREAD_ID-05:LSYS_ID-32:RT:Found: session id 0x12. sess
tok 28685

Nov  7 13:21:47 13:21:47.217761:CID-0:THREAD_ID-05:LSYS_ID-32:RT:  flow got session.

Nov  7 13:21:47 13:21:47.217761:CID-0:THREAD_ID-05:LSYS_ID-32:RT:  flow session id 18

Nov  7 13:21:47 13:21:47.217763:CID-0:THREAD_ID-05:LSYS_ID-32:RT: vector bits 0x200 vector
0x84ae85f0

Nov  7 13:21:47 13:21:47.217764:CID-0:THREAD_ID-05:LSYS_ID-32:RT:set nat 0x11e463550(18) timeout
const to 2

Nov  7 13:21:47 13:21:47.217765:CID-0:THREAD_ID-05:LSYS_ID-32:RT: set_nat_timeout 2 on session 18

Nov  7 13:21:47 13:21:47.217765:CID-0:THREAD_ID-05:LSYS_ID-32:RT:refresh nat 0x11e463550(18)
timeout to 2

Nov  7 13:21:47 13:21:47.217767:CID-0:THREAD_ID-05:LSYS_ID-32:RT:insert usp tag for apps

Nov  7 13:21:47 13:21:47.217768:CID-0:THREAD_ID-05:LSYS_ID-32:RT:mbuf 0x4882b600, exit nh
0xfffb0006
```

show log filename (QFabric System)

```
user@qfabric> show log messages
Mar 28 18:00:06 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:06 ED1486 chassisd:
CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on (jnxFruContentsIndex 8, jnxFruL1Index 1,
jnxFruL2Index 1, jnxFruL3Index 0, jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11,
jnxFruSlot 0, jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 2159)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1486 chassisd:
```

```
CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on (jnxFruContentsIndex 8, jnxFruL1Index 1,
jnxFruL2Index 2, jnxFruL3Index 0, jnxFruName PIC:  @ 0/1/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 2191)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492 chassisd:
CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on (jnxFruContentsIndex 8, jnxFruL1Index 1,
jnxFruL2Index 1, jnxFruL3Index 0, jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11,
jnxFruSlot 0, jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 242726)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492 chassisd:
CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on (jnxFruContentsIndex 8, jnxFruL1Index 1,
jnxFruL2Index 2, jnxFruL3Index 0, jnxFruName PIC:  @ 0/1/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 242757)
Mar 28 18:00:16 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:16 ED1486 file: UI_COMMIT:
User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:27 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:27 ED1486 file: UI_COMMIT:
User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50 _DCF_default___NW-
INE-0_RE0_ file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50 _DCF_default___NW-
INE-0_RE0_ file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:55 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:55 ED1492 file: UI_COMMIT:
User 'root' requested 'commit' operation (comment: none)
Mar 28 18:01:10 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:01:10 ED1492 file: UI_COMMIT:
User 'root' requested 'commit' operation (comment: none)
Mar 28 18:02:37 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:02:37 ED1491 chassisd:
CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on (jnxFruContentsIndex 8, jnxFruL1Index 1,
jnxFruL2Index 1, jnxFruL3Index 0, jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11,
jnxFruSlot 0, jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 33809)
```

show log user

```
user@host> show log user
usera    mg2546                  Thu Oct  1 19:37   still logged in
usera    mg2529                  Thu Oct  1 19:08 - 19:36  (00:28)
usera    mg2518                  Thu Oct  1 18:53 - 18:58  (00:04)
root     mg1575                  Wed Sep 30 18:39 - 18:41  (00:02)
root     ttyp2    aaa.bbbb.com   Wed Sep 30 18:39 - 18:41  (00:02)
userb    ttyp1    192.0.2.0      Wed Sep 30 01:03 - 01:22  (00:19)
```

show log accepted-traffic (SRX4600, SRX5400, SRX5600, and SRX5800)

```
user@host> show log accepted-traffic
Jul 17 20:26:04  sourpunch RT_FLOW: RT_FLOW_SESSION_CREATE: session created 3.3.3.5/2-
>4.4.4.2/63 0x0 None 3.3.3.5/2->4.4.4.2/63 0x0 N/A N/A N/A N/A 17 p2 TRUST UNTRUST 2617282058
N/A(N/A) xe-7/0/0.0 UNKNOWN UNKNOWN UNKNOWN N/A N/A -1 N/A N/A N/A
Jul 17 20:26:04  sourpunch RT_FLOW: RT_FLOW_SESSION_CREATE: session created 3.3.3.4/4-
>4.4.4.2/63 0x0 None 3.3.3.4/4->4.4.4.2/63 0x0 N/A N/A N/A N/A 17 p2 TRUST UNTRUST 2550162754
N/A(N/A) xe-7/0/0.0 UNKNOWN UNKNOWN UNKNOWN N/A N/A -1 N/A N/A N/A
Jul 17 20:26:04  sourpunch RT_FLOW: RT_FLOW_SESSION_CREATE: session created 3.3.3.4/1-
>4.4.4.2/63 0x0 None 3.3.3.4/1->4.4.4.2/63 0x0 N/A N/A N/A N/A 17 p2 TRUST UNTRUST 2550162755
N/A(N/A) xe-7/0/0.0 UNKNOWN UNKNOWN UNKNOWN N/A N/A -1 N/A N/A N/A
Jul 17 20:26:04  sourpunch RT_FLOW: RT_FLOW_SESSION_CREATE: session created 3.3.3.3/0-
>4.4.4.2/63 0x0 None 3.3.3.3/0->4.4.4.2/63 0x0 N/A N/A N/A N/A 17 p2 TRUST UNTRUST 2550162752
N/A(N/A) xe-7/0/0.0 UNKNOWN UNKNOWN UNKNOWN N/A N/A -1 N/A N/A N/A
Jul 17 20:26:04  sourpunch RT_FLOW: RT_FLOW_SESSION_CREATE: session created 3.3.3.5/5-
>4.4.4.2/63 0x0 None 3.3.3.5/5->4.4.4.2/63 0x0 N/A N/A N/A N/A 17 p2 TRUST UNTRUST 2550162751
N/A(N/A) xe-7/0/0.0 UNKNOWN UNKNOWN UNKNOWN N/A N/A -1 N/A N/A N/A
Jul 17 20:26:04  sourpunch RT_FLOW: RT_FLOW_SESSION_CREATE: session created 3.3.3.3/3-
>4.4.4.2/63 0x0 None 3.3.3.3/3->4.4.4.2/63 0x0 N/A N/A N/A N/A 17 p2 TRUST UNTRUST 2550162753
N/A(N/A) xe-7/0/0.0 UNKNOWN UNKNOWN UNKNOWN N/A N/A -1 N/A N/A N/A
```

Release Information

Command introduced before Junos OS Release 7.4.

Option *device-type (device-id | device-alias)* is introduced in Junos OS Release 13.1 for the QFX Series.

RELATED DOCUMENTATION

- syslog (System)

show route tenant

IN THIS SECTION

- Syntax | 974
- Description | 974
- Options | 975
- Required Privilege Level | 975
- Output Fields | 975
- Sample Output | 976
- Release Information | 976

Syntax

```
show route tenant
< (all | tenant-name)>
```

Description

Displays the routing table information for the tenant systems.

A tenant system is a logical partitioning of the device into multiple security domains similar to logical systems, to provide high scalability. A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. A set of interfaces that belong to the routing instances and the routing protocol parameters control the information in the routing tables. A tenant system can configure the assigned routing instances and the interfaces within a tenant system.

Options

all Displays the summary statistics about the all entries in the routing table.

tenant-name Specify the tenant system name.

Required Privilege Level

view

Output Fields

Table 43 on page 975 lists the output fields for the `show route tenant` command. Output fields are listed in the approximate order in which they appear.

Table 43: show route summary Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.o).
destinations	Number of destinations that correspond to the routes in the routing table.
routes	Number of routes in the routing table: • **active**—Number of routes that are active. • **holddown**—Number of routes that are in the hold-down state before being declared inactive. • **hidden**—Number of routes that are not used because of routing policy.
Direct	Routes on the directly connected network.
Local	Specify the local routes.

Table 43: show route summary Output Fields *(Continued)*

Field Name	Field Description
OSPF	Name of the protocol from which the route was learnt.

Sample Output

show route tenant tenant1

```
user@host> show route tenant tenant1
tenant1.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
12.1.1.0/24        *[Direct/0] 00:33:07
> via ge-0/0/0.0
12.1.1.1/32        *[Local/0] 00:33:07
Local via ge-0/0/0.0
224.0.0.5/32       *[OSPF/10] 00:30:31, metric 1
MultiRecv
tenant1.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
ff02::2/128        *[INET6/0] 00:33:33
MultiRecv
```

Release Information

Statement introduced in Junos OS Release 18.3R1.

show security application-firewall rule-set

IN THIS SECTION

- Syntax | 977
- Description | 977
- Options | 978
- Required Privilege Level | 978
- Output Fields | 978
- Sample Output | 980
- Sample Output | 980
- Sample Output | 981
- Release Information | 982

Syntax

```
show security application-firewall rule-set (<rule-set-name> | all)
show security application-firewall rule-set (rule-set-name | all) | (logical-system logical-system-name  | all) | all-logical-systems-tenants | root-logical-system | tenant (tenant-name  | all)
```

Description

Display information about the specified rule set defined in the application firewall.

The application firewall is defined by a collection of rule sets. A rule set defines the rules that specify match criteria, including dynamic applications, and the action to be taken for matching traffic.

Starting in Junos OS Release 18.2R1, the application firewall (AppFW) functionality is deprecated. As a part of this change, the `[edit security application-firewall]` hierarchy and all the configuration options under this hierarchy are deprecated— rather than immediately removed—to provide backward

compatibility and an opportunity to bring your configuration into compliance with the new configuration.

Options

rule-set-name	Display the name of the rule set.
all	(default) Display all rule sets for all logical systems. The user logical system administrator can display all rule sets only for the logical system they can access.
logical-system-name	Display application firewall rule set information for a specific logical system.
root-logical-system	Display application firewall rule set information for the root logical system (primary administrator only).
all-logical-systems-tenants	Display application firewall rule set information for all the logical systems and tenants.
tenant	Display application firewall rule set information for the tenant systems.

Required Privilege Level

view

Output Fields

Table 44 on page 978 lists the output fields for the `show security application-firewall rule-set` command. Output fields are listed in the approximate order in which they appear.

Table 44: show security application-firewall rule-set Output Fields

Field Name	Field Description
Rule-set	Name of the rule set.

Table 44: show security application-firewall rule-set Output Fields *(Continued)*

Field Name	Field Description
Logical system	Name of the logical system of the rule set.
Tenant	Name of the tenant system of the rule set.
Profile	The redirect profile to be used for rules requiring redirection for reject or deny actions.

Rule	Name of the rule
	• **Dynamic applications**—Name of the applications.
	• **Dynamic application groups**—Name of the application groups.
	• **SSL-Encryption**—Setting for SSL traffic.
	• **Action**—The action taken with respect to a packet that matches the application firewall rule set. Actions include the following:
	  • **permit**
	  • **deny**
	  • **reject**
	  • **redirect**
	• **Number of sessions matched**—Number of sessions matched with the application firewall rule.
	• **Number of sessions redirected**—Number of sessions redirected by the application firewall rule.

Default rule	The default rule applied when the identified application is not specified in any rules of the rule set. • **Number of sessions matched**—Number of sessions matched with the application firewall default rule. • **Number of sessions redirected**—Number of sessions redirected by the application firewall rule.

Table 44: show security application-firewall rule-set Output Fields *(Continued)*

Field Name	Field Description
Number of sessions with appid pending	Number of sessions that are pending application identification processing

Sample Output

show security application-firewall rule-set my_ruleset1

```
user@host>show security application-firewall rule-set my_ruleset1
Rule-set: my_ruleset1
     Rule: rule1
         Dynamic Applications: junos:FACEBOOK-ACCESS, junos:YMSG
         Dynamic Application Groups: junos:web, junos:chat
         SSL-Encryption: any
         Action: deny or redirect
         Number of sessions matched: 10
         Number of sessions redirected: 10
     Default rule: permit
         Number of sessions matched: 200
         Number of sessions redirected: 0
     Number of sessions with appid pending: 2
```

Sample Output

show security application-firewall rule-set all

```
user@host> show security application-firewall rule-set all

Rule-set: ls-product-design-rs1
    Logical system: ls-product-design
    Rule: r1
        Dynamic Applications: junos:TELNET
```

```
        Action:permit
        Number of sessions matched: 10
Default rule:deny
        Number of sessions matched: 100
Number of sessions with appid pending: 2

Rule-set: ls-product-design-rs1
    Logical system: ls-product-design
    Rule: r2
        Dynamic Application Groups: junos:web
        Action:permit
        Number of sessions matched: 20
Default rule:deny
        Number of sessions matched: 200
Number of sessions with appid pending: 4

Rule-set: ls-product-design-rs2
    Logical system: ls-product-design
    Rule: r1
        Dynamic Applications: junos:FACEBOOK-ACCESS
        Action:deny
        Number of sessions matched: 40
Default rule:permit
        Number of sessions matched: 400
Number of sessions with appid pending: 10
```

Sample Output

show security application-firewall rule-set ruleset1 tenant all

```
user@host> show security application-firewall rule-set ruleset1 tenant all

Rule-set: ruleset1
    Logical system: root-logical-system
    Tenant: TSYS1
    Rule: rule1
        Dynamic Applications: junos:HTTP, junos:FTP
        SSL-Encryption: any
        Action:permit
```

```
        Number of sessions matched: 0
        Number of sessions redirected: 0
Default rule:permit
        Number of sessions matched: 0
        Number of sessions redirected: 0
Number of sessions with appid pending: 0
```

Release Information

Command introduced in Junos OS Release 11.1. Updated in Junos OS Release 12.1X44-D10 with output format changes. Updated in Junos OS Release 12.1X45-D10 with redirection counters.

The `tenant` and `all-logical-systems-tenants` options are introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

clear security application-firewall rule-set statistics

show security application-firewall rule-set logical-system

IN THIS SECTION

- Syntax | **983**
- Description | **983**
- Options | **983**
- Required Privilege Level | **984**
- Output Fields | **984**
- Sample Output | **985**
- Release Information | **987**

Syntax

The primary, or root, administrator can issue the following statements:

```
show security application-firewall rule-set all
show security application-firewall rule-set rule-set-name | all | logical-system logical-system-name | all | root-logical-system [logical-system-name | all ]
```

The user logical system administrator can issue the following statement:

```
show security application-firewall rule-set all
```

Description

Display information about application firewall rule set(s) associated with a specific logical system, all logical systems, or the root logical system configured on a device.

NOTE: The primary administrator can configure and view application firewall rule sets for the root logical system and all user logical systems configured on the device. User logical system administrators can configure and view application firewall rule set information only for the user logical systems for which they have access. For information about primary and user administrator roles in logical systems, see Understanding Logical Systems for SRX Series Services Gateways.

Starting in Junos OS Release 18.2R1, the application firewall (AppFW) functionality is deprecated. As a part of this change, the `[edit security application-firewall]` hierarchy and all the configuration options under this hierarchy are deprecated— rather than immediately removed—to provide backward compatibility and an opportunity to bring your configuration into compliance with the new configuration.

Options

rule-set-name—Name of a specific rule set.

logical-system-name—Name of a specific logical system.

all—(default) Display all rule sets for all logical systems. The user logical system administrator can display all rule sets only for the logical system they can access.

root-logical-system—Display application firewall rule set information for the root logical system (primary administrator only).

Required Privilege Level

view

Output Fields

Table 45 on page 984 lists the output fields for the show security application-firewall rule-set logical-system command. Output fields are listed in the approximate order in which they appear.

Table 45: show security application-firewall rule-set logical-system Output Fields

Field Name	Field Description
Rule-set	Name of the rule set.
Logical system	Name of the logical system.

Rule	Name of the rule.
	• Dynamic applications—Name of the applications.
	• Dynamic application groups—Name of the application groups.
	• Action—The action taken with respect to a packet that matches the application firewall rule set. Actions include the following:
	  • permit
	  • deny
	• Number of sessions matched—Number of sessions matched with the application firewall rule.

Table 45: show security application-firewall rule-set logical-system Output Fields *(Continued)*

Field Name	Field Description
`Default rule`	The default rule applied when the identified application is not specified in any rules of the rule set. • `Number of sessions matched`—Number of sessions matched with the application firewall default rule.
`Number of sessions with appid pending`	Number of sessions that are pending with the application ID processing.

Sample Output

show security application-firewall rule-set logical-system all

```
root@host> show security application-firewall rule-set logical-system all

Rule-set: root_rs1
    Logical system: root-logical-system
    Rule: r1
        Dynamic Applications: junos:FTP
        Action:permit
        Number of sessions matched: 10
Default rule:deny
        Number of sessions matched: 100
Number of sessions with appid pending: 4

Rule-set: root-rs2
    Logical system: root-logical-system
    Rule: r1
        Dynamic Application Groups: junos:web
        Action:permit
        Number of sessions matched: 20
Default rule:deny
        Number of sessions matched: 100
```

```
Number of sessions with appid pending: 10
```

show security application-firewall rule-set all

```
root@host> show security application-firewall rule-set all

Rule-set: ls-product-design-rs1
    Logical system: ls-product-design
    Rule: r1
        Dynamic Applications: junos:TELNET
        Action:permit
        Number of sessions matched: 10
Default rule:deny
        Number of sessions matched: 100
Number of sessions with appid pending: 2

Rule-set: ls-product-design-rs1
    Logical system: ls-product-design
    Rule: r2
        Dynamic Application Groups: junos:web
        Action:permit
        Number of sessions matched: 20
Default rule:deny
        Number of sessions matched: 200
Number of sessions with appid pending: 4

Rule-set: ls-product-design-rs2
    Logical system: ls-product-design
    Rule: r1
        Dynamic Applications: junos:FACEBOOK-ACCESS
        Action:deny
        Number of sessions matched: 40
Default rule:permit
        Number of sessions matched: 400
Number of sessions with appid pending: 10
```

Release Information

Command introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

clear security application-firewall rule-set statistics logical-system

show security application-tracking counters

IN THIS SECTION

- Syntax | **987**
- Description | **987**
- Required Privilege Level | **988**
- Output Fields | **988**
- Sample Output | **988**
- Release Information | **989**

Syntax

```
show security application-tracking counters
```

Description

Display the status of AppTrack counters. These counters provide number of times an Apptrack message —that is—AppTrack session create, session close, route changes, and volume are generated.

Required Privilege Level

view

Output Fields

Table 46 on page 988 lists the output fields for the `show security application-tracking counters` command. Output fields are listed in the approximate order in which they appear.

Table 46: show security application-tracking counters

Field Name	Field Description
`Session create messages`	The number of log messages generated when a session was created.
`Session close messages`	The number of log messages generated when a session was closed.
`Session volume updates`	The number of log messages generated when an update interval was exceeded.
`Session route updates`	The number of log messages generated when an egress interface was selected based on application carried in the session by APBR.
`Failed messages`	The number of messages that were not generated due to memory or session constraints.

Sample Output

show security application-tracking counters

```
user@host> show security application-tracking counters

Application tracking counters:
```

```
    AppTrack counter type                    Value
Session create messages                        1
Session close messages                         1
Session volume updates                         0
Session route updates                          1
Failed messages                                0
```

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Understanding AppTrack

Example: Configuring AppTrack

show security alg status logical-system

IN THIS SECTION

- Syntax | **990**
- Description | **990**
- Options | **990**
- Additional Information | **990**
- Required Privilege Level | **990**
- Output Fields | **991**
- Sample Output | **991**
- Release Information | **993**

Syntax

```
show security alg status logical-system
<logical-system-name>
```

Description

Display the ALG status for a specific logical system or for all logical systems on the device.

Options

logical-system-name	Display ALG status for specific logical system.
all	Display ALG status for all logical systems.

Additional Information

The `show security alg status` command is used to view the ALG status in root logical system. The `show security alg status logical-system lsys1` command is used to view the ALG status in logical system lsys1. The `show security alg status logical-system all` command is used to view the ALG status of all existing logical systems.

NOTE: Only users under root logical system can view the ALG status for all logical systems. The keyword logical-system is not required in the command `show security alg status logical-system lsys1` when you log in to a particular logical system.

Required Privilege Level

view

Output Fields

Sample Output

show security alg status logical-system all

```
user@host> show security alg status logical-system all
Logical system: root-logical-system
ALG Status:
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Disabled
  SQL      : Disabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled

Logical system: LSYS2
ALG Status:
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Disabled
  SQL      : Disabled
```

```
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled

Logical system: LSYS0
ALG Status:
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Disabled
  SQL      : Disabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled

Logical system: LSYS1
ALG Status:
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Disabled
  SQL      : Disabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled

{secondary:node0}
```

show security alg status logical-system LSYS1

```
user@host> show security alg status logical-system LSYS1
ALG Status:
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Disabled
  SQL      : Disabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled

{secondary:node0}
```

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

Understanding Application Layer Gateway (ALG) in Logical Systems | 305

alg

Example: Enabling FTP ALG in a Logical System | 311

show security datapath-debug capture

IN THIS SECTION

- Syntax | 994
- Description | 994
- Required Privilege Level | 994
- Output Fields | 994
- Sample Output | 995
- Release Information | 995

Syntax

```
show security datapath-debug capture
```

Description

Display details of the data path debugging capture file.

Required Privilege Level

view

Output Fields

Output fields are listed in the approximate order in which they appear.

Sample Output

show security datapath—debug capture

```
user@host> show security datapath-debug capture
Packet 1, len 120: (C0/F0/P0/SEQ:71:lbt)
91 00 00 47 11 00 10 00 9a 14 00 19 03 00 00 00
00 00 00 00 00 01 00 47 10 00 00 00 00 00 00 00
00 1f 12 f8 dd 29 00 21 59 84 f4 01 81 00 02 1e
08 00 45 60 01 f4 00 00 00 00 3f 06 73 9f 01 01
01 02 03 01 01 02 d4 31 d4 31 00 00 00 00 00 00
00 00 50 02 00 00 ff ad 00 00 00 00
Packet 2, len 120: (C0/F0/P0/SEQ:71:lbt)
90 00 00 47 04 00 00 00 00 00 00 00 02 02 00 47
10 00 00 00 00 00 00 00 50 00 a6 1c 00 00 00 00
00 00 00 0a 00 00 00 00 00 00 09 d9 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 1f 12 f8
dd 29 00 21 59 84 f4 01 81 00 02 1e
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

show security datapath-debug counter | 996

Understanding Data Path Debugging for Logical Systems

show security datapath-debug counter

IN THIS SECTION

- Syntax | 996
- Description | 996
- Required Privilege Level | 996
- Output Fields | 996
- Sample Output | 997
- Release Information | 997

Syntax

```
show security datapath-debug counter
```

Description

Display details of the data path debugging counter.

Required Privilege Level

view

Output Fields

Output fields are listed in the approximate order in which they appear.

Sample Output

show security datapath-debug counter

```
user@host> show security datapath-debug counter
Datapath debug counters
Packet Filter 1:
np-ingress
Chassis 0 FPC 4 : 1
np-ingress
Chassis 0 FPC 3 : 0
np-egress
Chassis 0 FPC 4 : 1
np-egress
Chassis 0 FPC 3 : 0
jexec
Chassis 0 FPC 0 PIC 1: 0
jexec
Chassis 0 FPC 0 PIC 0: 1
lbt
Chassis 0 FPC 0 PIC 1: 0
lbt
Chassis 0 FPC 0 PIC 0: 2
pot
Chassis 0 FPC 0 PIC 1: 0
pot
```

Release Information

Command introduced in Junos OS Release 10.0.

RELATED DOCUMENTATION

show security datapath-debug capture | 994

Understanding Data Path Debugging for Logical Systems

show security dns-cache

IN THIS SECTION

- Syntax | 998
- Description | 998
- Options | 999
- Required Privilege Level | 999
- Output Fields | 999
- Sample Output | 999
- Sample Output | 1000
- Release Information | 1000

Syntax

```
show security dns-cache  <dns-name dns-name>
```

Description

Display DNS cache information.

NOTE: This command is only available to the primary administrator on devices that are configured for logical systems. This command is not available in user logical systems or on devices that are not configured for logical systems.

Options

- dns-name—Display DNS cache information for the specified name.

Required Privilege Level

view

Output Fields

Table 47 on page 999 lists the output fields for the show security dns-cache command. Output fields are listed in the approximate order in which they appear.

Table 47: show security dns-cache Output Fields

Field Name	Field Description
DNS Name	DNS name.
Address Family	IPv4 or IPv6.
TTL	Time-to-live value.
IP Address	IP address for the DNS name.

Sample Output

show security dns-cache

```
user@host> show security dns-cache
DNS Name: dns1.test.com:
  Address Family: IPv4, TTL: 10
```

```
    IP Address: 10.1.1.1
  Address Family: IPv6: TTL = 15
    IP Address: 2001:10.1.1.1
DNS Name: dns2.test.com:
  Address Family: IPv4, TTL: 20
    IP Address: 10.2.2.2
    IP Address: 10.2.2.3
```

Sample Output

show security dns-cache dns-name dns2.test.com

```
user@host> show security dns-cache dns-name dns2.test.com
DNS Name: dns2.test.com:
  Address Family: IPv4, TTL: 20
    IP Address: 10.2.2.2
    IP Address: 10.2.2.3
```

Release Information

Command introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

clear security dns-cache | 904

show security dynamic-address

IN THIS SECTION

- Syntax | 1001
- Description | 1001
- Options | 1002
- Required Privilege Level | 1003
- Output Fields | 1004
- Sample Output | 1005
- Release Information | 1011

Syntax

```
show security dynamic-address
<summary>
<category-name (Blacklist | CC | GeoIP | IPFilter | Infected-Hosts | SecProfiling | Whitelist) >
<family inet | inet6>
<feed-name (feed-name) >
<address-name (address-name) >
<ip-start (starting-IP-address) >
<ip-end (ending-IP-address) >
<instance (advanced-anti-malware | default | geoip) >
<logical-system (logical-system-name | all)>
<tenant (tenant-name | all)>
```

Description

Displays information about dynamic addresses. Each dynamic address belongs to only one instance. Within that instance is a set of categories to which the dynamic address further belongs.

A dynamic address entry provides dynamic IP address information to security policies. A dynamic address entry is a group of IP addresses, not just a single IP prefix, that can be imported in from external sources. These IP addresses are for specific domains or for entities that have a common attribute such as a particular undesired location that poses a threat. The administrator can then configure security policies to use the DAE within a security policy.

Options

none Display source category (feed) and dynamic address name for all nodes (primary and backup nodes in case of HA.) The same as the `show security dynamic-address ip-start 1.0.0.0 ip-end 255.255.255.255` command.

summary (Optional) Display basic information of dynamic-address including their name, feeds, properties and number of IPv4 and IPv6 entries.

category-name (*category-name*) (Optional) Display the source category (feed) and dynamic address name for the specified threat type (category name). A category is basically a list of feeds of the same type. The type defines SRX Series enforcement point criteria for feed lookup and enforcement. Supported category names are:

- Blocklist—A list of locations (IP addresses, URLs, etc.) that you do not trust. A blocklist allows everyone access except for the members on the blocklist.
- CC—A list of known C&C servers that are able to send commands to members of a botnet.
- GeoIP—A list giving you the ability to filter traffic to and from specific geographies in the world.
- IPFilter—A list of addresses and ranges of malicious sites that can send junk data.
- Infected-Hosts—A lists of hosts within your network that may have been compromised and require attention. This list is generated from Juniper ATP Cloud based on hosts that have downloaded malware.
- SecProfiling—A security intelligence category. You configure the security policy or IDP rule to add the source IP or destination IP of the traffic to the secProfiling feed in the cloud. You can also configure the security policy or IDP to match the source IP or destination IP of the traffic against the secProfiling feed in the cloud.
- Allowlist—A list of locations (IP addresses, URLs, etc.) that you trust. An allowlist denies everyone access except for the members on the allowlist.

family	(Optional) Show the dynamic-address for specified protocol-family. Both IPv4 and IPv6 are displayed if no family is specified. 'inet' and 'inet6' can be combined with other options of the show command. For example, show security dynamic-address family inet6 ip-start 1111::1 ip-end 3333::3.
feed-name	(Optional) User-defined name of the source feed.
address-name	(Optional) The dynamic address name. If you do not specify an `address-name`, then information related to all dynamic addresses downloaded to this SRX Series device is displayed.
ip-start	(Optional) The numerical minimum IP address where you want to investigate. Specifying `ip-start` and optionally `ip-end` (it is not required to specify `ip-end` if you use `ip-start`) is helpful to filter the output to a specific range instead of having to review the entire list which can be very long.
ip-end	(Optional) The numerical maximum IP address where you want to investigate. Specifying `ip-start` and optionally `ip-end` is helpful to filter the output to a specific range instead of having to review the entire list which can be very long. If you specify `ip-end`, you must specify `ip-start`.
instance (*instance-name*)	(Optional) The physically separated database. Supported instance names are: • advanced-anti-malware—The IP-based allowlists and blocklists. • default—The default instance holds the following data: blocklist, allowlist, C&C, infected host, and IPfilter. • geoip—The geoip data.
logical-system (logical-system-name | all)	Perform this operation on all logical systems or on a particular logical system.
tenant (tenant-name | all)	Perform this operation on all tenant systems or on a particular tenant system.

Required Privilege Level

View

Output Fields

Table 48 on page 1004 describes the output fields for the `show security dynamic-address` command. Output fields are listed in the approximate order in which they appear.

Table 48: show security dynamic-address Output Fields

Field Name	Field Description
Address name	Dynamic address entry name.
Address ID	Internal ID used to uniquely identify the dynamic address entry.
IPv4 entries	The number of IPv4 entries in the specific dynamic address.
IPv6 entries	The number of IPv6 entries in the specific dynamic address.
Category/feed	The threat type associated with the dynamic address. See description of `category-name` and `feed-name` above.
Total number of IPv4 entries	The number of IPv4 entries in the database.
Total number of IPv4 entries from feed	The number of IPv4 entries in the feed. An entry in a feed can correspond to multiple entries in the database.
Total number of IPv6 entries	The number of IPv6 entries in the database.
Total number of IPv6 entries from feed	The number of IPv6 entries in the feed. An entry in a feed can correspond to multiple entries in the database.
Instance default	Total number of default matching entries.
Instance geoip	Total number of geoip data matching entries.
Instance advanced-anti-malware	Total number of the IP-based allowlists and blocklists matching entries.

Table 48: show security dynamic-address Output Fields *(Continued)*

Field Name	Field Description
Session-scan	Status of the session scan.

Sample Output

show security dynamic-address summary

```
user@host> show security dynamic-address summary
Dynamic-address session scan status            : Disable
Hold-interval for dynamic-address session scan : 10 seconds




  Server Name                 : fd_svr1
    Hostname/IP               : https://192.208.132.20
    Update interval           : 1000
    Hold   interval           : 1200
    TLS Profile Name          : tls_profile_1
    User        Name          : ---




    Feed Name                             : fd1
        Mapped dynamic address name       : ad1
        URL                               : https://192.208.132.20/fd1.gz
        Feed update interval              : 1000     Feed hold interval :1200
        Total update                      : 3
        Total IPv4 entries                : 1
        Total IPv6 entries                : 1
        Total download   errors           : 0        Last occurence N/A
        Total db         errors           : 0        Last occurence N/A
        Total other      errors           : 0        Last occurence N/A
        Total ageout                      : 0        Last occurence N/A
```

```
        Next update time                 : Mon Aug 31 16:19:44 2020
        Next expire time                 : Mon Aug 31 16:53:04 2020
        Flags                            : 0x0
        Last update file size            : 47
        Last update IPv4 entries         : 1
        Last update IPv6 entries         : 1
        Last update   begin time         : Mon Aug 31 16:03:04 2020
        Last update   end   time         : Mon Aug 31 16:03:04 2020
        Last update   time cost(s)       : --
        Last download begin time         : Mon Aug 31 16:03:04 2020
        Last download end   time         : Mon Aug 31 16:03:04 2020
        Last update   status             : 255
        Last download time cost(s)       : --
    Feed Name                            : fd2
        Mapped dynamic address name      : ad2
        URL                              : https://192.208.132.20/fd2.gz
        Feed update interval             : 1000     Feed hold interval :1200
        Total update                     : 2
        Total IPv4 entries               : 1
        Total IPv6 entries               : 1
        Total download   errors          : 0        Last occurence N/A
        Total db         errors          : 0        Last occurence N/A
        Total other      errors          : 0        Last occurence N/A
        Total ageout                     : 0        Last occurence N/A
        Next update time                 : Mon Aug 31 16:19:44 2020
        Next expire time                 : Mon Aug 31 16:53:04 2020
        Flags                            : 0x0
        Last update file size            : 47
        Last update IPv4 entries         : 1
        Last update IPv6 entries         : 1
        Last update   begin time         : Mon Aug 31 16:03:04 2020
        Last update   end   time         : Mon Aug 31 16:03:04 2020
        Last update   time cost(s)       : --
        Last download begin time         : Mon Aug 31 16:03:04 2020
        Last download end   time         : Mon Aug 31 16:03:04 2020
        Last update   status             : 255
        Last download time cost(s)       : --



    Address name                : ad1
      Address id                : 11
```

```
      IPv4 entries              : 1
      IPv6 entries              : 1
      Session-scan              : Disable
      Mapping Feed        : fd1
      Category/feed       : ---     / ---
            property name : feed
                    value : fd1




    Address name                : ad2
      Address id                : 12
      IPv4 entries              : 1
      IPv6 entries              : 1
      Session-scan              : Disable
      Mapping Feed        : fd2
      Category/feed       : ---     / ---
            property name : feed
                    value : fd2




Instance Name                             : default
    Total number of IPv4 entries          : 3
    Total number of IPv4 entries from feed : 2
    Total number of IPv6 entries          : 3
    Total number of IPv6 entries from feed : 2



Instance Name                             : geoip
    Total number of IPv4 entries          : 0
    Total number of IPv6 entries          : 0



Instance Name                             : advanced-anti-malware
    Total number of IPv4 entries          : 0
    Total number of IPv6 entries          : 0
```

show security dynamic-address summary

```
user@host> show security dynamic-address summary
node1:
--------------------------------------------------------------------------


    Address name          : a1
    Address id            : 11
      IPv4 entries        : 13778
      IPv6 entries        : 0
      Category/feed       : GeoIP   / ---
            property name : country
                    value : AU
                    value : CN


    Address name          : a2
    Address id            : 12
      IPv4 entries        : 0
      IPv6 entries        : 0
      Category/feed       : IPFilter / ---
            property name : test
                    value : test


Total number of IPv4 entries           : 13778
Total number of IPv4 entries from feed : 0
Total number of IPv4 except entries    : 0
Total number of IPv6 entries           : 0
Total number of IPv6 entries from feed : 0
```

show security dynamic-address instance advanced-anti-malware

```
user@host> show security dynamic-address instance advanced-anti-malware
node1:
--------------------------------------------------------------------------
No.      IP-start        IP-end          Feed             Address
1        5.5.0.0         5.5.0.10        global_whitelist ID-00000003
2        11.11.0.0       11.11.0.10      global_blacklist ID-00000004
```

show security dynamic-address instance geoip

```
user@host> show security dynamic-address instance geoip
node1:
--------------------------------------------------------------------------
No.      IP-start        IP-end          Feed             Address
1        1.0.0.0         1.0.0.255       geoip_country    a1
2        1.0.1.0         1.0.1.255       geoip_country    a1
3        1.0.2.0         1.0.3.255       geoip_country    a1
4        1.0.4.0         1.0.7.255       geoip_country    a1
5        1.0.8.0         1.0.15.255      geoip_country    a1
6        1.0.32.0        1.0.63.255      geoip_country    a1
7        1.1.0.0         1.1.0.255       geoip_country    a1
8        1.1.1.0         1.1.1.255       geoip_country    a1
9        1.1.2.0         1.1.3.255       geoip_country    a1
10       1.1.4.0         1.1.7.255       geoip_country    a1
11       1.1.8.0         1.1.15.255      geoip_country    a1
12       1.1.16.0        1.1.31.255      geoip_country    a1
13       1.1.32.0        1.1.63.255      geoip_country    a1
14       1.2.0.0         1.2.1.255       geoip_country    a1
```

show security dynamic-address category-name Infected-Hosts

```
user@host> show security dynamic-address category-name Infected-Hosts

node1:
--------------------------------------------------------------------------
No.          IP-start            IP-end          Feed                      Address
1               1.0.0.7                 1.0.0.7              Infected-Hosts/1
ID-21500011
2             1.0.0.10             1.0.0.10         Infected-Hosts/1         ID-21500011
3             1.0.0.21             1.0.0.21         Infected-Hosts/1         ID-21500011
4             1.0.0.11             1.0.0.11         Infected-Hosts/1         ID-21500012
5             1.0.0.12             1.0.0.12         Infected-Hosts/1         ID-21500012
6             1.0.0.22             1.0.0.22         Infected-Hosts/1         ID-21500012
7             1.0.0.6                 1.0.0.6             Infected-Hosts/1         ID-21500013
8             1.0.0.9                 1.0.0.9             Infected-Hosts/1         ID-21500013
9             1.0.0.13             1.0.0.13         Infected-Hosts/1         ID-21500013
10             1.0.0.23             1.0.0.23         Infected-Hosts/1         ID-21500013
```

show security dynamic-address logical-system LSYS1

```
user@host> show security dynamic-address logical-system LSYS1

Instance default Total number of matching entries: 0
Instance geoip   Total number of matching entries: 0
Instance advanced-anti-malware Total number of matching entries: 0
```

show security dynamic-address logical-system all

```
user@host> show security dynamic-address logical-system all

Instance default Total number of matching entries: 0
Instance geoip   Total number of matching entries: 0
Instance advanced-anti-malware Total number of matching entries: 0
```

show security dynamic-address tenant TSYS1

```
user@host> show security dynamic-address tenant TSYS1

Instance default Total number of matching entries: 0
Instance geoip   Total number of matching entries: 0
Instance advanced-anti-malware Total number of matching entries: 0
```

show security dynamic-address tenant all

```
user@host> show security dynamic-address tenant all

Instance default Total number of matching entries: 0
Instance geoip   Total number of matching entries: 0
Instance advanced-anti-malware Total number of matching entries: 0
```

Release Information

Command introduced in Junos OS Release 12.1X46-D25.

Command <family inet | inet6> introduced in Junos OS Release 18.1.

The logical-system and tenant options are introduced in Junos OS Release 18.4R1.

show security firewall-authentication history

IN THIS SECTION

- Syntax | 1011
- Description | 1012
- Options | 1012
- Required Privilege Level | 1012
- Output Fields | 1013
- Sample Output | 1013
- Sample Output | 1014
- Release Information | 1014

Syntax

```
show security firewall-authentication history
<address (address)>
<from-zone (from-zone)>
<identifier (identifier)>
<logical-system (logical-system-name | all)>
<node (node-id | all | local | primary)>
<root-logical-system (address | from-zone | identifier | tenant | to-zone)>
<tenant (tenant-name | all)>
<to-zone (to-zone)>
```

Description

Displays security firewall authentication user history information and verify the number of firewall users who successfully authenticated and the number of firewall users who failed to log in.

Options

- none—Display history of firewall authentication information.
- address—Display authentication entries based on IP address.
- from-zone—Display authentication entries matching the given source zone, null for web-authentication and userfw-authentication.
- identifier—Display authentication entries by user identifier.
- logical-system—Display firewall authentication tables based on logical system name.
- `node`—(Optional) For chassis cluster configurations, display all firewall authentication history on a specific node (device) in the cluster.
  - *`node-id`* —Identification number of the node. It can be 0 or 1.
  - `all`—Display information about all nodes.
  - `local`—Display information about the local node.
  - `primary`—Display information about the primary node.
- root-logical-system—Display firewall authentication tables for root logical system.
- tenant—Display firewall authentication tables based on tenant name.
- to-zone—Display authentication entry matching the given destination zone, null for web-auth and userfw-auth.

Required Privilege Level

view

Output Fields

Table 49 on page 1013 lists the output fields for the `show security firewall-authentication history` command. Output fields are listed in the approximate order in which they appear.

Table 49: show security firewall-authentication history Output Fields

Field Name	Field Description
`Authentications`	Number of authentications.
`Id`	Identification number.
`Source IP`	IP address of the authentication source.
`Date`	Authentication date.
`Time`	Authentication time.
`Duration`	Authentication duration.
`Status`	Authentication status success or failure.
`User`	Name of the user.

Sample Output

show security firewall-authentication history

```
user@host> show security firewall-authentication history
History of firewall authentication data:
  Authentications: 1
```

```
          Id Source Ip       Date       Time     Duration   Status   User
           1 203.0.113.1     2007-04-03 11:43:06 00:00:45   Success  hello
```

Sample Output

show security firewall-authentication history node all

```
user@host> show security firewall-authentication history node all
node0:
--------------------------------------------------------------------------
History of firewall authentication data:
Authentications: 2
Id Source Ip       Date       Time     Duration   Status   User
1 203.0.113.1      2008-01-04 12:00:10 0:05:49    Success  local1
2 203.0.113.1      2008-01-04 14:36:52 0:01:03    Success  local1
node1:
--------------------------------------------------------------------------
History of firewall authentication data:
  Authentications: 1
          Id Source Ip       Date       Time     Duration   Status   User
           203.0.113.1       2008-01-04 14:59:43 1193046:06: Success local1
```

show security firewall-authentication history tenant tn1

```
user@host> show security firewall-authentication history tenant tn1
History of firewall authentication data:
  Authentications: 0
```

Release Information

Command introduced in Junos OS Release 8.5. The `node` option is added in Junos OS Release 9.0. The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Understanding Logical System Firewall Authentication | 185

Firewall User Authentication Overview

show security firewall-authentication users

IN THIS SECTION

- Syntax | 1015
- Description | 1016
- Options | 1016
- Required Privilege Level | 1016
- Output Fields | 1017
- Sample Output | 1017
- Sample Output | 1018
- Sample Output | 1018
- show security firewall-authentication users tenant all | 1019
- Release Information | 1019

Syntax

```
show security firewall-authentication users
<address (ip-address )>
<auth-type (pass-through | user-firewall | web-authentication)>
<from-zone (from-zone)>
<identifier (identifier )>
<logical-system (logical-system-name | all)>
<node (node-id | all | local | primary)>
<root-logical-system (address | auth-type | from-zone | identifier | tenant | to-zone>
<tenant (tenant-name |all)>
<to-zone (to-zone )>
```

Description

Display firewall authentication details about all users and verify the number of firewall users who successfully authenticated and firewall users who failed to log in.

Options

- none—Display details about all firewall authentication users.
- address—Display authentication entries based on ip address.
- auth-type—Display authentication entries matching the given auth-type.
- from-zone—Display authentication entries matching the given source zone, null for web-auth and userfw-auth.
- identifier—Display authentication entries by id.
- logical-system—Display firewall authentication tables based on logical system name.
- `node`—(Optional) For chassis cluster configurations, display firewall authentication details for all users on a specific node.
  - *node-id*—Identification number of the node. It can be 0 or 1.
  - `all`—Display information about all nodes.
  - `local`—Display information about the local node.
  - `primary`—Display information about the primary node.
- root-logical-system—Display firewall authentication tables for root logical system.
- tenant—Display firewall authentication tables based on tenant name.
- to-zone—Display authentication entry matching the given destination zone, null for web-auth and userfw-auth.

Required Privilege Level

view

Output Fields

Table 50 on page 1017 lists the output fields for the `show security firewall-authentication users` command. Output fields are listed in the approximate order in which they appear.

Table 50: show security firewall-authentication users Output Fields

Field Name	Field Description
`Total users in table`	Gives count of how many entries/users the command will display.
`Id`	Identification number.
`Source IP`	IP address of the authentication source.
`Src zone`	User traffic received from the zone.
`Dst zone`	User traffic destined to the zone.
`Profile`	Name of profile used for authentication.
`Age`	Idle timeout for the user.
`Status`	Authentication status success or failure.
`User`	Name of the user.

Sample Output

show security firewall-authentication users

```
user@host> show security firewall-authentication users
Firewall authentication data:
  Total users in table: 1
```

```
          Id Source Ip       Src zone Dst zone Profile    Age Status   User
           1 192.0.2.5/24         z1       z2       p1       0 Success  local1
```

Sample Output

show security firewall-authentication users node 0

```
user@host> show security firewall-authentication users node 0
node0:
--------------------------------------------------------------------------
Firewall authentication data:
  Total users in table: 1
          Id Source Ip       Src zone Dst zone Profile    Age Status   User
           3 192.0.2.5/24        z1       z2       p1           1 Success  local1
```

Sample Output

show security firewall-authentication users node all

```
user@host> show security firewall-authentication users node all
node0:
--------------------------------------------------------------------------
Firewall authentication data:
  Total users in table: 1
          Id Source Ip       Src zone Dst zone Profile    Age Status   User
           3 192.0.2.5        z1       z2       p1           1 Success  local1
node1:
--------------------------------------------------------------------------
Firewall authentication data:
  Total users in table: 1
          Id Source Ip       Src zone Dst zone Profile    Age Status   User
           2 192.0.2.5       z1       z2       p1           1 Success  local1
```

show security firewall-authentication users tenant all

command-name

```
user@host> show security firewall-authentication users tenant all
Firewall authentication data:
  Total users in table: 1
          Id Source Ip                         Src zone Dst zone Profile    Age Status
User
           2 192.0.2.10                       N/A      N/A      test-rad     1 Success
b1
```

Release Information

Command introduced in Junos OS Release 8.5. The `node` options added in Junos OS Release 9.0. The `tenant` option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Firewall User Authentication Overview

show security flow session

IN THIS SECTION

- Syntax | **1020**
- Description | **1020**
- Options | **1020**
- Required Privilege Level | **1023**
- Output Fields | **1023**
- Sample Output | **1027**

- Release Information | **1033**

Syntax

```
show security flow session [filter] [brief | extensive | summary]
<node (node-id  | all | local | primary)>
```

Description

Display information about all currently active security sessions on the device. For the normal flow sessions, the `show security flow session` command displays byte counters based on IP header length. However, for sessions in Express Path mode, the statistics are collected from the IOC2 (SRX5K-MPC), IOC3 (SRX5K-MPC3-100G10G and SRX5K-MPC3-40G10G), and IOC4 (SRX5K-IOC4-MRAT and SRX5K-IOC4-10G) ASIC hardware engines and include full packet length with L2 headers. Because of this, the output displays slightly larger byte counters for sessions in Express Path mode than for the normal flow session.

Options

- *filter*—Filter the display by the specified criteria.

  The following filters reduce the display to those sessions that match the criteria specified by the filter. Refer to the specific `show` command for examples of the filtered output.

  advanced-anti-malware — Show advanced-anti-malware sessions. For details on the `advanced-anti-malware` option, see the Sky Advanced Threat Prevention CLI Reference Guide.

  all-logical-systems-tenants — All multitenancy systems.

  application — Predefined application name.

  application-firewall — Application firewall enabled.

application-firewall-rule-set	Application firewall enabled with the specified rule set.
application-traffic-control	Application traffic control session.
application-traffic-control-rule-set	Application traffic control rule set name and rule name.
bytes-less-than	Define session's bytes-count less than a value (1..4294967295).
bytes-more-than	Define session's bytes-count more a value (1..4294967295).
conn-tag	Session connection tag (0..4294967295).
content-filtering	Display the content filtering session details.
curr-less-than	Define session's current-timeout value less than a value (1..100000).
curr-more-than	Define session's current-timeout value more than a value (1..100000).
destination-port	Destination port.
destination-prefix	Destination IP prefix or address.
dynamic-application	Dynamic application.
dynamic-application-group	Dynamic application.
duration-less-than	Define session's duration time less than a value (1..100000).
duration-more-than	Define session's duration time more than a value (1..100000).
encrypted	Encrypted traffic.
family	Display session by family.
ha-link	Display HA link session information.
idp	IDP-enabled sessions.
interface	Name of incoming or outgoing interface.
logical-system (all | *logical-system-name*)	Name of a specific logical system or `all` to display all logical systems.
nat	Display sessions with network address translation.
node	(Optional) For chassis cluster configurations, display security flow session information on a specific node (device) in the cluster.

- *node-id* —Identification number of the node. It can be 0 or 1.
- all —Display information about all nodes.
- local —Display information about the local node.
- primary—Display information about the primary node.

packets-less-than	Define session's packets-count less than a value (1..4294967295).
packets-more-than	Define session's packets-count more than a value (1..4294967295).
plugin-name	Plugin name.
plugin-status	Plugin status.
plugins	Display the flow session information of plugins.
policy-id	Display session information based on policy ID; the range is 1 through 4,294,967,295.
pretty	Display the flow session information in a list to make it easy for you to read and monitor.
protocol	IP protocol number.
resource-manager	Resource manager.
root-logical-system	Display root logical system as default.
security-intelligence	Display security intelligence sessions.
services-offload	Display services offload sessions.
session-identifier	Display session with specified session identifier.
session-state	Session state.
source-port	Source port.
source-prefix	Source IP prefix.
ssl	Display the SSL proxy sessions information.
tenant	Displays the security flow session information for a tenant system.
timeout-less-than	Define session's timeout value less than a value (1..100000).
timeout-more-than	Define session's timeout value more than a value (1..100000).
tunnel	Tunnel sessions.

tunnel-inspection-type	Tunnel inspection type	
	gre	Displays gre tunnel inspection
	ipip	Displays ipip tunnel inspection
	vxlan	Displays vxlan tunnel inspection
vxlan-vni	It only lists the tunnel session which vni matches the one you specify in the command.	
url-category	Display flow session information by url-category.	
vrf-group	Display flow session information by L3VPN VRF Group.	
web-filtering	Display the web filtering sessions details.	

- brief | extensive | summary—Display the specified level of output.
- none—Display information about all active sessions.

Required Privilege Level

view

Output Fields

Table 51 on page 1023 lists the output fields for the `show security flow session` command. Output fields are listed in the approximate order in which they appear.

Table 51: show security flow session Output Fields

Field Name	Field Description	Level of Output
`Session ID`	Number that identifies the session. Use this ID to get more information about the session.	brief extensive none

Table 51: show security flow session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
If	Interface name.	brief none
State	Status of security flow session.	brief extensive none
Conn Tag	A 32-bit connection tag that uniquely identifies the GPRS tunneling protocol, user plane (GTP-U) and the Stream Control Transmission Protocol (STCP) sessions. The connection tag for GTP-U is the tunnel endpoint identifier (TEID) and for SCTP is the vTag. The connection ID remains 0 if the connection tag is not used by the sessions.	brief extensive none
CP Session ID	Number that identifies the central point session. Use this ID to get more information about the central point session.	brief extensive none
Policy name	Name and ID of the policy that the first packet of the session matched.	brief extensive none
Timeout	Idle timeout after which the session expires.	brief extensive none

Table 51: show security flow session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`In`	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).	brief extensive none
Bytes	Number of received and transmitted bytes.	brief extensive none
Pkts	Number of received and transmitted packets.	brief extensive none
`Total sessions`	Total number of sessions.	brief extensive none
`Out`	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).	brief extensive none
`Status`	Session status.	extensive
`Flag`	Internal flag depicting the state of the session, used for debugging purposes.	extensive
`Source NAT pool`	The name of the source pool where NAT is used.	extensive
`Dynamic application`	Name of the application.	extensive

Table 51: show security flow session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Application traffic control rule-set`	AppQoS rule set for this session.	extensive
`Rule`	AppQoS rule for this session.	extensive
`Maximum timeout`	Maximum session timeout.	extensive
`Current timeout`	Remaining time for the session unless traffic exists in the session.	extensive
`Session State`	Session state.	extensive
`Start time`	Time when the session was created, offset from the system start time.	extensive
`Unicast-sessions`	Number of unicast sessions.	Summary
`Multicast-sessions`	Number of multicast sessions.	Summary
`Services-offload-sessions`	Number of services-offload sessions.	Summary
`Failed-sessions`	Number of failed sessions.	Summary
`Sessions-in-use`	Number of sessions in use. • Valid sessions • Pending sessions • Invalidated sessions • Sessions in other states	Summary

Table 51: show security flow session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Maximum-sessions	Maximum number of sessions permitted.	Summary

Sample Output

show security flow session

```
root> show security flow session
Flow Sessions on FPC0 PIC1:

Session ID: 10115977, Policy name: SG/4, State: Active, Timeout: 56, Valid
  In: 203.0.113.1/1000 --> 203.0.113.11/2000;udp, Conn Tag: 0x0, If: reth1.0, Pkts: 1, Bytes:
86, CP Session ID: 10320276
  Out: 203.0.113.11/2000 --> 203.0.113.1/1000;udp, Conn Tag: 0x0, If: reth0.0, Pkts: 0, Bytes:
0, CP Session ID: 10320276

Total sessions: 1
```

show security flow session (with default policy)

```
root> show security flow session
Session ID: 36, Policy name: pre-id-default-policy/n, Timeout: 2, Valid
  In: 10.10.10.2/61606 --> 10.10.10.1/179;tcp, Conn Tag: 0x0, If: ge-0/0/2.0, Pkts: 1, Bytes: 64,
  Out: 10.10.10.1/179 --> 10.10.10.2/61606;tcp, Conn Tag: 0x0, If: .local..0, Pkts: 1, Bytes: 40,
```

show security flow session (drop flow)

Shows dropped flows for SRX5400.

```
root> show security flow session
Outgoing wing:   CP session ID: 12, CP sess SPU Id: 4617   1.0.0.1/55069 <- 1.0.0.254/23;6,
Conn, Drop Flow
```

```
Tag: 0x0, VRF GRP ID: 0(0), If: xe-1/0/0.0 (7), Flag: 0x40000020, Vector index: 0x00000002
WSF: 1, Diff: 0, Sequence: 0, Ack: 0, Port sequence: 0, FIN sequence: 0, FIN state: 0   Zone Id:
7, NH: 0x40010, NSP tunnel: 0x0, NP info: 0xffthread id:255
```

show security flow session (IPv6 tunnel)

```
root> show security flow session
Session ID: 44, Policy name: N/A, State: Stand-alone, Timeout: N/A, Valid
  In: 9001::4/1 --> 9001::3/1;ipip, Conn Tag: 0x0, If: ge-0/0/0.0, Pkts: 0, Bytes: 0,

Session ID: 45, Policy name: N/A, State: Stand-alone, Timeout: N/A, Valid
  In: 9001::4/1 --> 9001::3/1;ipv6, Conn Tag: 0x0, If: ge-0/0/0.0, Pkts: 0, Bytes: 0,

Session ID: 57, Policy name: default-policy-logical-system-00/2, State: Stand-alone, Timeout:
1796, Valid
  In: 20.0.0.2/37628 --> 30.0.0.2/22;tcp, Conn Tag: 0x0, If: ge-0/0/1.0, Pkts: 22, Bytes: 4409,
  Out: 30.0.0.2/22 --> 20.0.0.2/37628;tcp, Conn Tag: 0x0, If: ip-0/0/0.1, Pkts: 30, Bytes: 5209,

Session ID: 58, Policy name: default-policy-logical-system-00/2, State: Stand-alone, Timeout:
1784, Valid
  In: 2001::2/58602 --> 3001::2/22;tcp, Conn Tag: 0x0, If: ge-0/0/1.0, Pkts: 31, Bytes: 5569,
  Out: 3001::2/22 --> 2001::2/58602;tcp, Conn Tag: 0x0, If: ip-0/0/0.1, Pkts: 28, Bytes: 6249,
Total sessions: 4
```

show security flow session brief

```
root> show security flow session brief
Flow Sessions on FPC0 PIC1:

Session ID: 10115977, Policy name: SG/4, State: Active, Timeout: 62, Valid
  In: 203.0.113.11/1000 --> 203.0.113.1/2000;udp, Conn Tag: 0x0, If: reth1.0, Pkts: 1, Bytes:
86, CP Session ID: 10320276
  Out: 203.0.113.1/2000 --> 203.0.113.11/1000;udp, Conn Tag: 0x0, If: reth0.0, Pkts: 0, Bytes:
0, CP Session ID: 10320276

Total sessions: 1
```

show security flow session content-filtering

```
root> show security flow session content-filtering
Flow Sessions on FPC0 PIC1:

Session ID: 10115977, Policy name: SG/4, State: Active, Timeout: 62, Valid
  In: 192.0.2.0/24/1000 --> 203.0.113.0/2000;udp, Conn Tag: 0x0, If: reth1.0, Pkts: 1, Bytes:
86, CP Session ID: 10320276
  Out: 203.0.113.0/2000 --> 192.0.2.0/24/1000;udp, Conn Tag: 0x0, If: reth0.0, Pkts: 0, Bytes:
0, CP Session ID: 10320276

Total sessions: 1
```

show security flow session extensive

```
root> show security flow session extensive
Flow Sessions on FPC0 PIC1:

Session ID: 10115977, Status: Normal, State: Active
Flags: 0x8000040/0x18000000/0x12000003
Policy name: SG/4
Source NAT pool: Null, Application: junos-gprs-gtp-v0-udp/76
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 90, Current timeout: 54
Session State: Valid
Start time: 6704, Duration: 35
   In: 203.0.113.11/1000 --> 201.11.0.100/2000;udp,
    Conn Tag: 0x0, Interface: reth1.0,
    Session token: 0x6, Flag: 0x40000021
    Route: 0x86053c2, Gateway: 201.10.0.100, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 86
    CP Session ID: 10320276
   Out: 203.0.113.1/2000 --> 203.0.113.11/1000;udp,
    Conn Tag: 0x0, Interface: reth0.0,
    Session token: 0x7, Flag: 0x50000000
    Route: 0x86143c2, Gateway: 203.0.113.11, Tunnel: 0
```

```
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
    CP Session ID: 10320276
Total sessions: 1
```

show security flow session extensive

```
root> show security flow session extensive
Flow Sessions on FPC0 PIC0:

Session ID: 10000059, Status: Normal
Flags: 0x10000/0x0/0x10/0x1
Policy name: N/A
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: N/A, Current timeout: N/A
Session State: Valid
Start time: 642, Duration: 369
   In: 3.0.0.2/64387 --> 2.0.0.1/8940;esp,
  Conn Tag: 0x0, Interface: xe-2/0/2.0,
    Session token: 0x7, Flag: 0x80100621
    Route: 0xc0010, Gateway: 2.0.0.2, Tunnel: 0
    ESP/AH frag Rx: 0, Generated: 0
    Inner IPv4 frag Rx: 0, Tx: 0, Generated: 0,
    Inner IPv6 frag Rx: 0, Tx: 0, Generated: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 25, Bytes: 3760
    CP Session ID: 0

Session ID: 10000060, Status: Normal
Flags: 0x10000/0x0/0x10/0x1
Policy name: N/A
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: N/A, Current timeout: N/A
```

```
Session State: Valid
Start time: 642, Duration: 369
   In: 3.0.0.2/0 --> 2.0.0.1/0;esp,
  Conn Tag: 0x0, Interface: xe-2/0/2.0,
    Session token: 0x7, Flag: 0x621
    Route: 0xc0010, Gateway: 2.0.0.2, Tunnel: 0
    ESP/AH frag Rx: 0, Generated: 0
    Inner IPv4 frag Rx: 0, Tx: 0, Generated: 0,
    Inner IPv6 frag Rx: 0, Tx: 0, Generated: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
    CP Session ID: 0
Total sessions: 2
```

show security flow session summary

```
root> show security flow session summary
Flow Sessions on FPC10 PIC1:
Unicast-sessions: 1
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 1
  Valid sessions: 1
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC10 PIC2:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456
```

```
Flow Sessions on FPC10 PIC3:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456
```

show security flow session tunnel-inspection-type

```
root> show security flow session tunnel-inspection-type vxlan
Session ID: 335544369, Policy name: p1/7, Timeout: 2, Valid
In: 192.168.200.100/19183 --> 192.168.200.101/2;icmp, Conn Tag: 0xfcd, If: xe-7/0/0.0, Pkts: 2,
Bytes: 2048, CP Session ID: 30, Tunnel Session ID: 268435486, Type: VXLAN, VNI: 1000
Out: 192.168.200.101/2 --> 192.168.200.100/19183;icmp, Conn Tag: 0xfcd, If: xe-7/0/1.0, Pkts: 2,
Bytes: 2048, CP Session ID: 30, Tunnel Session ID: 268435488, Type: VXLAN, VNI: 1000
```

show security flow session tunnel-inspection-type

```
root> show security flow session vxlan-vni 400
Session ID: 1677861258, Policy name: pset1_p1/6, Timeout: 2, Valid
In: 192.150.0.12/55908 --> 192.160.0.66/80;tcp, Conn Tag: 0xfcd, If: xe-3/0/0.0, Pkts: 5, Bytes:
465, CP Session ID: 7021087,
Type: VXLAN, VNI: 400, Tunnel Session ID: 1680264845
Out: 192.160.0.66/80 --> 192.150.0.12/55908;tcp, Conn Tag: 0xfcd, If: xe-3/0/1.0, Pkts: 3,
Bytes: 328, CP Session ID: 7021087,
Type: VXLAN, VNI: 400, Tunnel Session ID: 1679640460



Session ID: 1678454648, Policy name: pset1_p1/6, Timeout: 2, Valid
In: 192.150.0.13/56659 --> 192.160.0.67/80;tcp, Conn Tag: 0xfcd, If: xe-3/0/0.0, Pkts: 5, Bytes:
465, CP Session ID: 5589311,
Type: VXLAN, VNI: 400, Tunnel Session ID: 1679698941
Out: 192.160.0.67/80 --> 192.150.0.13/56659;tcp, Conn Tag: 0xfcd, If: xe-3/0/1.0, Pkts: 3,
```

```
Bytes: 328, CP Session ID: 5589311,
Type: VXLAN, VNI: 400, Tunnel Session ID: 1679872223
```

show security flow session web-filtering

```
root> show security flow session web-filtering
Session ID: 256, Policy name: p/4, Timeout: 1794, Session State: Valid
  In: 198.51.100.0/33170 --> 203.0.113.0/443;tcp, Conn Tag: 0x0, If: ge-0/0/0.0, Pkts: 3, Bytes:
351,
  Out: 203.0.113.0/443 --> 192.0.2.0/13089;tcp, Conn Tag: 0x0, If: ge-0/0/2.0, Pkts: 1, Bytes:
44,
Total sessions: 1
```

Release Information

Command introduced in Junos OS Release 8.5.

Support for filter and view options added in Junos OS Release 10.2.

Application firewall, dynamic application, and logical system filters added in Junos OS Release 11.2.

Policy ID filter added in Junos OS Release 12.3X48-D10.

Support for connection tag added in Junos OS Release 15.1X49-D40.

The `tenant` option introduced in Junos OS Release 18.3R1.

The `tunnel-inspection-type` option is introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

clear security flow session all

show security flow session tenant

IN THIS SECTION

- Syntax | 1034
- Description | 1034
- Options | 1034
- Required Privilege Level | 1035
- Output Fields | 1035
- Sample Output | 1036
- Release Information | 1037

Syntax

```
show security flow session tenant (tenant-name | all)
```

Description

Displays the information about the currently active security flow sessions of the tenant systems on the device. You can either view the currently active security flow sessions information for a specific tenant system or for all the tenant systems.

Options

tenant-name	Name of the tenant system.
all	Displays the security flow session information for all the tenant systems.

Required Privilege Level

view

Output Fields

Table 52 on page 1035 lists the output fields for the `show security flow session tenant` command. Output fields are listed in the approximate order in which they appear.

Table 52: show security flow session tenant

Field Name	Field Description
`Session ID`	Number that identifies the session. You can use this ID to get additional information about the session.
`Policy name`	Policy that permitted the traffic.
`Timeout`	Idle timeout after which the session expires.
`In`	Incoming security flow session details. The incoming security flow session details include the source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets, and bytes.
`Out`	Reverse security flow session details. The reverse security flow session details include the source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets, and bytes.
`Total sessions`	Total number of security flow sessions.
`Policy name`	Name and ID of the policy that the first packet of the security flow session matched.
`Tenant`	Name of the tenant system.

Sample Output

show security flow session tenant T1

```
root@host> show security flow session tenant T1
Flow Sessions on FPC7 PIC1:

Session ID: 290000224, Policy name: default-policy-logical-system-32/2, Timeout: 1790, Valid
Tenant: T1
  In: 203.0.113.0/39767 --> 203.0.113.1/23;tcp, Conn Tag: 0x0, If: xe-3/0/1.0, Pkts: 39, Bytes:
2136, CP Session ID: 1225556754
  Out: 203.0.113.1/23 --> 203.0.113.0/39767;tcp, Conn Tag: 0x0, If: lt-0/0/0.101, Pkts: 31,
Bytes: 1872, CP Session ID: 1225556754
```

show security flow session tenant all

```
root@host> show security flow session tenant all
Flow Sessions on FPC7 PIC1:

Session ID: 290000224, Policy name: default-policy-logical-system-32/2, Timeout: 1790, Valid
Tenant: T1
  In: 203.0.113.0/39767 --> 203.0.113.1/23;tcp, Conn Tag: 0x0, If: xe-3/0/1.0, Pkts: 39, Bytes:
2136, CP Session ID: 1225556754
  Out: 203.0.113.1/23 --> 203.0.113.0/39767;tcp, Conn Tag: 0x0, If: lt-0/0/0.101, Pkts: 31,
Bytes: 1872, CP Session ID: 1225556754

Session ID: 290000225, Policy name: default-policy-logical-system-33/2, Timeout: 1790, Valid
Tenant: T2
  In: 203.0.113.3/39767 --> 203.0.113.4/23;tcp, Conn Tag: 0x0, If: lt-0/0/0.103, Pkts: 39,
Bytes: 2136, CP Session ID: 1225556755
  Out: 203.0.113.4/23 --> 203.0.113.3/39767;tcp, Conn Tag: 0x0, If: xe-9/0/0.0, Pkts: 31, Bytes:
1872, CP Session ID: 1225556755
Total sessions: 2

Flow Sessions on FPC7 PIC2:
Total sessions: 0
```

```
Flow Sessions on FPC7 PIC3:
Total sessions: 0
```

Release Information

Command introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

clear security flow session tenant | 940

show security idp logical system

IN THIS SECTION

- Syntax | 1037
- Description | 1038
- Options | 1038
- Required Privilege Level | 1038
- Sample Output | 1038
- Release Information | 1038

Syntax

```
show security idp  logical-system
```

Description

Display information about the logical systems and the IDP policies associated to the logical systems.

Options

- `logical-system`— Show the IDP policy assigned to a logical system. The IDP policy is assigned to a logical system through the security profile.
  - `policy-association`

Required Privilege Level

view

Sample Output

show security idp logical-system policy-association

```
user@host> show security idp logical-system policy-association
Logical system                  IDP policy

root-logical-system             policy1
LSYS1                           idpengine
```

Release Information

Command introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

clear security idp

show security idp attack table

IN THIS SECTION

- Syntax | 1039
- Description | 1039
- Options | 1040
- Required Privilege Level | 1040
- Output Fields | 1040
- Sample Output | 1041
- Release Information | 1042

Syntax

```
show security idp attack table
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays the detailed information of IDP attack table and displays the IPv6 covert channels which are identified and mitigated.

Options

none	Displays the details of the IDP attack table.
logical-system *logical-system-name*	(Optional) Displays the details of the IDP attack table for a specific logical system.
logical-system all	(Optional) Displays the details of the IDP attack table for all logical systems.
tenant *tenant-name*	(Optional) Displays the details of the IDP attack table for a specific tenant system.

Required Privilege Level

view

Output Fields

Table 53 on page 1040 lists the output fields for the `show security idp attack table` command. Output fields are listed in the approximate order in which they appear.

Table 53: show security idp attack table Output Fields

Field Name	Field Description
`Attack name`	Name of the attack that you want to match in the monitored network traffic.

Table 53: show security idp attack table Output Fields *(Continued)*

Field Name	Field Description
Hits	Total number of attack matches. On SRX Series devices, for brute force and time-binding-related attacks, the logging is to be done only when the match count is equal to the threshold. That is, only one log is generated within the 60-second period in which the threshold is measured. This process prevents repetitive logs from being generated and ensures consistency with other IDP platforms, such as IDP-standalone. When no attack is seen within the 60-second period and the BFQ entry is flushed out, the match count starts over the new attack match shows up in the attack table, and the log is generated.

Sample Output

show security idp attack table

```
user@host> show security idp attack table
IDP attack statistics:
  Attack name                                #Hits
  HTTP:OVERFLOW:PI3WEB-SLASH-OF              1
```

show security idp attack table tenant TSYS1

```
user@host> show security idp attack table tenant TSYS1
IDP attack statistics:

  Attack name                                #Hits
  FTP:USER:ROOT                              1
```

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

IPv6 covert channels are detected in Junos OS Release 19.1R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

clear security idp attack table

show security idp counters action

IN THIS SECTION

- Syntax | **1042**
- Description | **1043**
- Options | **1043**
- Required Privilege Level | **1043**
- Output Fields | **1043**
- Sample Output | **1044**
- Release Information | **1045**

Syntax

```
show security idp counters action
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays the detailed information of IDP counter type and value.

Options

none	Displays the detailed information of IDP counter type and value.
logical system *logical-system-name*	(Optional) Displays the detailed information of IDP counter type and value for a specific logical system.
logical system all	(Optional) Displays the detailed information of IDP counter type and value for all logical systems.
tenant *tenant-name*	(Optional) Displays the detailed information of IDP counter type and value for a specific tenant system.

Required Privilege Level

view

Output Fields

Table 54 on page 1043 lists the output fields for the `show security idp counters action` command. Output fields are listed in the approximate order in which they appear.

Table 54: show security idp counters action Output Fields

Field Name	Field Description
IDP counter type	Name of the action
Value	Number of packets dropped, recommended, and ignored based on the action in the IDP counters.

Sample Output

show security idp counters action

```
user@host> show security idp counters action
IDP counters:

  IDP counter type                                                  Value
 None                                                                0
 Recommended                                                         0
 Ignore                                                              0
 Diffserv                                                            0
 Drop packet                                                         0
 Drop                                                                0
 Close                                                               0
 Close server                                                        0
 Close client                                                        0
 IP action rate limit                                                0
 IP action drop                                                      0
 IP action close                                                     0
 IP action nofity                                                    0
 IP action failed                                                    0
```

show security idp counters action logical-system LSYS0

```
user@host> show security idp counters action logical-system LSYS0
IDP counters:

  IDP counter type                                                  Value
 None                                                                0
 Recommended                                                         0
 Ignore                                                              0
 Diffserv                                                            0
 Drop packet                                                         0
 Drop                                                                0
 Close                                                               0
 Close server                                                        0
 Close client                                                        0
 IP action rate limit                                                0
 IP action drop                                                      0
```

```
 IP action close                                                  0
 IP action nofity                                                 0
 IP action failed                                                 0
```

show security idp counters action tenant TSYS1

```
user@host> show security idp counters action tenant TSYS1
IDP counters:

  IDP counter type                                              Value
 None                                                             1
 Recommended                                                      0
 Ignore                                                           0
 Diffserv                                                         0
 Drop packet                                                      0
 Drop                                                             0
 Close                                                            0
 Close server                                                     0
 Close client                                                     0
 IP action rate limit                                             0
 IP action drop                                                   0
 IP action close                                                  0
 IP action nofity                                                 0
 IP action failed                                                 0
```

Release Information

Command introduced for user logical system in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

clear security idp counters action | 926

show security idp counters application-identification

IN THIS SECTION

- Syntax | 1046
- Description | 1046
- Options | 1046
- Required Privilege Level | 1047
- Output Fields | 1047
- Sample Output | 1050
- Release Information | 1052

Syntax

```
show security idp counters application-identification
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays the status of all IDP application identification (AI) counter values.

Options

none	Displays the status of all IDP application identification (AI) counter values.
logical-system *logical-system-name*	(Optional) Displays the status of all IDP application identification (AI) counter values for a specific logical system.

logical-system all	(Optional) Displays the status of all IDP application identification (AI) counter values for all logical systems.
tenant *tenant-name*	(Optional) Displays the status of all IDP application identification (AI) counter values for a specific tenant system.

Required Privilege Level

view

Output Fields

Table 55 on page 1047 lists the output fields for the `show security idp counters application-identification` command. Output fields are listed in the approximate order in which they appear.

Table 55: show security idp counters application-identification Output Fields

Field Name	Field Description
`AI matches`	Number of sessions with an AI signature match.
`AI no-matches`	Number of sessions with no AI signature match.
`AI-enabled sessions`	Number of sessions with AI enabled.
`AI-disabled sessions`	Number of sessions with AI disabled.
`AI-disabled sessions due to ssl encapsulated flows`	Number of sessions with AI disabled due to SSL encapsulated flows.
`AI-disabled sessions due to cache hit`	Number of sessions with AI disabled due to a cache match.

Table 55: show security idp counters application-identification Output Fields *(Continued)*

Field Name	Field Description
AI-disabled sessions due to configuration	Number of sessions with AI disabled because the configured session limit was reached.
AI-disabled sessions due to protocol remapping	Number of sessions with AI disabled due to protocol remapping.
AI-disabled sessions due to RPC match	Number of sessions with AI disabled due to an RPC match.
AI-disabled sessions due to gate match	Number of sessions with AI disabled due to a gate match.
AI-disabled sessions due to non-TCP/UDP flows	Number of sessions with AI disabled due to non-TCP or non-UDP flows.
AI-disabled sessions due to session limit	Number of sessions with AI disabled because the maximum session limit was reached.
AI-disabled sessions due to session packet memory limit	Number of sessions with AI disabled because the memory usage limit per session was reached.
AI-disabled sessions due to global packet memory limit	Number of sessions with AI disabled because the global memory usage limit was reached.
AI sessions current global reass packet memory usage	Number of AI sessions with current global reassembler packet memory usage limit
AI sessions peak global reass packet memory usage	Number of AI sessions with peak global reassembler packet memory usage limit
AI sessions current global packet memory usage	Number of AI sessions with current global packet memory usage limit

Table 55: show security idp counters application-identification Output Fields *(Continued)*

Field Name	Field Description
AI sessions peak global packet memory usage	Number of AI sessions with peak global packet memory usage limit
AI-sessions dropped due to malloc failure before session create	Number of AI sessions dropped because the malloc failure occurred before session create.
AI-sessions dropped due to malloc failure after create	Number of AI sessions dropped because the malloc failure occurred after session create.
AI-Packets received on sessions marked for drop due to malloc failure	Number of AI packets received on sessions that are marked to be dropped because the malloc failure.
Packets cloned for AI	Number of packets cloned for application identification.
Policy update	Number of times the IDP policy has been updated.
Total PME prematch job ignored	Number of jobs ignored because of pattern matching engine (PME) not matching.
Total packets for which prematch job were ignored	Number of packets for which signature matching was ignored as prematch found.
Prematch busy packet count	Number of packets saved as they are handed off for signature matching during prematch reprocess.
Final match busy packet count	Number of packets saved as they are handed off for signature matching during final match reprocess.
Total AI busy packet count	Number of times AI saved packet handed off for signature matching.

Table 55: show security idp counters application-identification Output Fields *(Continued)*

Field Name	Field Description
Final match processed busy packet count	Number of times a packet processed for final matching before signature matching.
Prematch processed busy packet count	Number of times a packet processed for prematch before signature match.
Prematch ignored busy packet count	Number of packets ignored for signature matching as prematch found.
AI done busy packet count	Number of packets signature matching not completed before AI done.
JPME flow for Ignored jobs destroyed	Number of jobs destroyed because of flow mismatch due to policy relookup.
Set AI done for prematch	Number of sessions set for AI applied.
AI done for prematch	Number of sessions with AI applied.

Sample Output

show security idp counters application-identification

```
user@host> show security idp counters application-identification

IDP counter type                                                Value
 AI matches                                                       0
 AI no-matches                                                    0
 AI-enabled sessions                                              0
 AI-disabled sessions                                             0
 AI-disabled sessions due to ssl encapsulated flows               0
 AI-disabled sessions due to cache hit                            0
```

```
 AI-disabled sessions due to configuration                                0
 AI-disabled sessions due to protocol remapping                           0
 AI-disabled sessions due to RPC match                                    0
 AI-disabled sessions due to gate match                                   0
 AI-disabled sessions due to non-TCP/UDP flows                            0
 AI-disabled sessions due to session limit                                0
 AI-disabled sessions due to session packet memory limit                  0
 AI-disabled sessions due to global packet memory limit                   0
 AI sessions current global reass packet memory usage                     0
 AI sessions peak global reass packet memory usage                        0
 AI sessions current global packet memory usage                           0
 AI sessions peak global packet memory usage                              0
 AI-sessions 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
 Packets cloned for AI                                                    0
 Policy update                                                            0
 Total PME prematch job ignored                                           0
 Total packets for which prematch job were ignored                        0
 Prematch busy packet count                                               0
 Final match busy packet count                                            0
 Total AI busy packet count                                               0
 Final match processed busy packet count                                  0
 Prematch processed busy packet count                                     0
 Prematch ignored busy packet count                                       0
 AI done busy packet count                                                0
 JPME flow for Ignored jobs destroyed                                     0
 Set AI done for prematch                                                 0
 AI done for prematch                                                     0
                                                                          0
```

show security idp counters application-identification tenant TSYS1

```
user@host> show security idp counters application-identification tenant TSYS1
IDP counters:

  IDP counter type                                                    Value
 AI matches                                                            0
 AI no-matches                                                         0
 AI-enabled sessions                                                   0
 AI-disabled sessions                                                  1
```

```
AI-disabled sessions due to ssl encapsulated flows                    0
AI-disabled sessions due to cache hit                                 1
AI-disabled sessions due to configuration                             0
AI-disabled sessions due to protocol remapping                        0
AI-disabled sessions due to RPC match                                 0
AI-disabled sessions due to gate match                                0
AI-disabled sessions due to non-TCP/UDP flows                         0
AI-disabled sessions due to global packet memory limit                0
AI sessions current global packet memory usage                        0
AI sessions peak global packet memory usage                           0
AI-sessions 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
Packets cloned for AI                                                 0
Policy update                                                         0
Total PME prematch job ignored                                        0
Total packets for which prematch job were ignored                     0
Prematch busy packet count                                            0
Final match busy packet count                                         0
Total AI busy packet count                                            0
Final match processed busy packet count                               0
Prematch processed busy packet count                                  0
Prematch ignored busy packet count                                    0
AI done busy packet count                                             0
JPME flow for Ignored jobs destroyed                                  0
Set AI done for prematch                                              0
AI done for prematch                                                  0
```

Release Information

Command introduced in Junos OS Release 9.2. Modified in Junos OS Release 12.1.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

clear security idp counters application-identification

show security idp counters memory

IN THIS SECTION

- Syntax | 1053
- Description | 1053
- Options | 1053
- Required Privilege Level | 1054
- Output Fields | 1054
- Sample Output | 1054
- Release Information | 1056

Syntax

```
show security idp counters memory
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays detailed information of allocated, reallocated IDP counters memory values.

Options

none	Displays detailed information of allocated, reallocated IDP counters memory values.
logical-system *logical-system-name*	(Optional) Displays detailed information of allocated, reallocated IDP counters memory values for a specific logical system.

logical-system all (Optional) Displays detailed information of allocated, reallocated IDP counters memory values for all logical systems.

tenant *tenant-name* (Optional) Displays detailed information of allocated, reallocated IDP counters memory values for a specific tenant system.

Required Privilege Level

view

Output Fields

Table 56 on page 1054 lists the output fields for the `show security idp counters memory` command. Output fields are listed in the approximate order in which they appear.

Table 56: show security idp counters memory Output Fields

Field Name	Field Description
IDP counter type	Name of the counter type
Value	Number of requests made for memory allocation and reallocation,

Sample Output

show security idp counters memory

```
user@host> show security idp counters memory
IDP counters:

  IDP counter type                                                Value
 Memory allocation requested                                      928058
 Memory reallocation requested                                    0
```

```
 Memory allocation failed                                    0
 Memory reallocation failed                                  0
 Memory free requested                                       889749
 Memory free failed                                          0
 IDP Arena requested                                         0
 IDP Arena failed                                            0
 IDP Arena freed                                             0
 Objcache requested                                          132
 Objcache failed                                             0
 Objcache freed                                              2
 Objcache over limit                                         0
 Objcache invalid record                                     0
 Detector Objcache requested                                 21032
 Detector Objcache failed                                    0
 Detector Objcache freed                                     20324
 Detector Objcache invalid record                            0
 Detector Arena requested                                    0
 Detector Arena failed                                       0
 Detector Arena freed                                        0
 Kzalloc requested                                           910823
 Kzalloc failed                                              0
 Kzalloc freed                                               885264
 Pool alloc called                                           0
 Pool alloc subscribed size                                  0
 Pool alloc done                                             0
 Pool free called                                            0
 Pool free done                                              0
 Malloc over limit                                           0
 Remote free deadbeef                                        0
```

show security idp counters memory tenant TSYS1

```
user@host> show security idp counters memory tenant TSYS1
IDP counters:

  IDP counter type                                          Value
 Memory allocation requested                                 78
 Memory reallocation requested                               0
 Memory allocation failed                                    0
 Memory reallocation failed                                  0
 Memory free requested                                       75
```

```
Memory free failed                                  0
IDP Arena requested                                 0
IDP Arena failed                                    0
IDP Arena freed                                     0
Objcache requested                                  70
Objcache failed                                     0
Objcache freed                                      68
Objcache over limit                                 0
Objcache invalid record                             0
Detector Objcache requested                         13
Detector Objcache failed                            0
Detector Objcache freed                             13
Detector Objcache invalid record                    0
Detector Arena requested                            0
Detector Arena failed                               0
Detector Arena freed                                0
Kzalloc requested                                   8
Kzalloc failed                                      0
Kzalloc freed                                       7
Pool alloc called                                   0
Pool alloc subscribed size                          0
Pool alloc done                                     0
Pool free called                                    0
Pool free done                                      0
Malloc over limit                                   0
Remote free deadbeef                                0
```

Release Information

Command introduced for user logical systems in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

clear security idp counters memory | 919

show security idp counters ssl-inspection

IN THIS SECTION

- Syntax | 1057
- Description | 1057
- Options | 1057
- Required Privilege Level | 1058
- Output Fields | 1058
- Sample Output | 1058
- Release Information | 1061

Syntax

```
show security idp counters ssl-inspection
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays the IDP counters value for decrypted and encrypted sessions.

Options

none	Displays the IDP counters value for decrypted and encrypted sessions.
logical-system *logical-system-name*	(Optional) Displays the IDP counters value for decrypted and encrypted sessions for a specific logical system.

logical-system all	(Optional) Displays the IDP counters value for decrypted and encrypted sessions for all logical systems.
tenant *tenant-name*	(Optional) Displays the IDP counters value for decrypted and encrypted sessions for a specific tenant system.

Required Privilege Level

view

Output Fields

Table 57 on page 1058 lists the output fields for the `show security idp counters ssl-inspection` command. Output fields are listed in the approximate order in which they appear.

Table 57: show security idp counters ssl-inspection Output Fields

Field Name	Field Description
IDP counter type	Name of the action
Value	Number of packets and sessions decrypted, sessions not decrypted.

Sample Output

show security idp counters ssl-inspection

```
user@host>  show security idp counters ssl-inspection
IDP counters:

  IDP counter type                                                Value
 Packets Decrypted                                                 0
 Sessions Decrypted                                                0
```

```
 Sessoins Not Decrypted                                          0
 Sessions Not Decrypted - Configuration                          0
 Sessions Not Decrypted - No Key                                 0
 Sessions Not Decrypted - Unsupported Ciphers                    0
 Sessions Not Decrypted - Unsupported Compression                0
 Sessions Not Decrypted - Unsupported Key Exchange               0
 Sessions Not Decrypted - Bulk Decryption Failure                0
 Sessions Not Decrypted - Key Generation Failure                 0
 Sessions Not Decrypted - Temporary Certificate                  0
 Sessions Not Decrypted - Handshake Verification Failure         0
 Sessions Not Decrypted - ID Cache Miss                          0
 Sessions Not Decrypted - Session Limit                          0
 Sessions Not Decrypted - Message Size                           0
 Sessions Not Decrypted - No Memory                              0
 Sessions New Key                                                0
 Sessions Used Key                                               0
 Session ID Cache Hits                                           0
 Session ID Cache Misses                                         0
 Sessions Used XLR RSA SAE for Key Decryption                    0
 Sessions - Error when XLR RSA SAE used                          0
```

show security idp counters ssl-inspection logical-system LSYS1

```
user@host>  show security idp counters ssl-inspection logical-system LSYS1
IDP counters:

  IDP counter type                                              Value
 Packets Decrypted                                               0
 Sessions Decrypted                                              0
 Sessoins Not Decrypted                                          0
 Sessions Not Decrypted - Configuration                          0
 Sessions Not Decrypted - No Key                                 0
 Sessions Not Decrypted - Unsupported Ciphers                    0
 Sessions Not Decrypted - Unsupported Compression                0
 Sessions Not Decrypted - Unsupported Key Exchange               0
 Sessions Not Decrypted - Bulk Decryption Failure                0
 Sessions Not Decrypted - Key Generation Failure                 0
 Sessions Not Decrypted - Temporary Certificate                  0
 Sessions Not Decrypted - Handshake Verification Failure         0
 Sessions Not Decrypted - ID Cache Miss                          0
 Sessions Not Decrypted - Session Limit                          0
```

```
 Sessions Not Decrypted - Message Size                                0
 Sessions Not Decrypted - No Memory                                   0
 Sessions New Key                                                     0
 Sessions Used Key                                                    0
 Session ID Cache Hits                                                0
 Session ID Cache Misses                                              0
 Sessions Used XLR RSA SAE for Key Decryption                         0
 Sessions - Error when XLR RSA SAE used                               0
```

show security idp counters ssl-inspection tenant TSYS1

```
user@host> show security idp counters ssl-inspection tenant TSYS1
IDP counters:

  IDP counter type                                                   Value
 Packets Decrypted                                                    0
 Sessions Decrypted                                                   0
 Sessoins Not Decrypted                                               0
 Sessions Not Decrypted - Configuration                               0
 Sessions Not Decrypted - No Key                                      0
 Sessions Not Decrypted - Unsupported Ciphers                         0
 Sessions Not Decrypted - Unsupported Compression                     0
 Sessions Not Decrypted - Unsupported Key Exchange                    0
 Sessions Not Decrypted - Bulk Decryption Failure                     0
 Sessions Not Decrypted - Key Generation Failure                      0
 Sessions Not Decrypted - Temporary Certificate                       0
 Sessions Not Decrypted - Handshake Verification Failure              0
 Sessions Not Decrypted - ID Cache Miss                               0
 Sessions Not Decrypted - Session Limit                               0
 Sessions Not Decrypted - Message Size                                0
 Sessions Not Decrypted - No Memory                                   0
 Sessions New Key                                                     0
 Sessions Used Key                                                    0
 Session ID Cache Hits                                                0
 Session ID Cache Misses                                              0
 Sessions Used XLR RSA SAE for Key Decryption                         0
 Sessions - Error when XLR RSA SAE used                               0
```

Release Information

Command introduced for user logical systems in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

clear security idp counters ssl-inspection | 917

show security idp counters pdf-decoder

IN THIS SECTION

- Syntax | 1061
- Description | 1062
- Options | 1062
- Required Privilege Level | 1062
- Output Fields | 1062
- Sample Output | 1063
- Release Information | 1065

Syntax

```
show security idp counters pdf-decoder
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays the IDP counters value for PDF-Decode enabled sessions, requests, and memory limit IDP counter types.

Options

none	Displays the IDP counters value for PDF-Decode enabled sessions, requests, and memory limit IDP counter types.
logical-system *logical-system-name*	(Optional) Displays the IDP counters value for PDF-Decode enabled sessions, requests, and memory limit IDP counter types for a specific logical system.
logical-system all	(Optional) Displays the IDP counters value for PDF-Decode enabled sessions, requests, and memory limit IDP counter types for all logical systems.
tenant *tenant-name*	(Optional) Displays the IDP counters value for PDF-Decode enabled sessions, requests, and memory limit IDP counter types for a specific tenant system.

Required Privilege Level

view

Output Fields

Table 58 on page 1062 lists the output fields for the `show security idp counters pdf-decoder` command. Output fields are listed in the approximate order in which they appear.

Table 58: show security idp counters pdf-decoder Output Fields

Field Name	Field Description
IDP counter type	Name of the action

Table 58: show security idp counters pdf-decoder Output Fields *(Continued)*

Field Name	Field Description
Value	Number of PDF-decode enabled sessions and requests for the IDP counters.

Sample Output

show security idp counters pdf-decoder

```
user@host>  show security idp counters pdf-decoder
IDP counters:

  IDP counter type                                                Value
 PDF-Decode enabled sessions                                       0
 PDF-Decode requests                                               0
 PDF-decode in pending state                                       0
 PDF-decode finished sucessfully                                   0
 PDF-decode per file memory limit reached                          0
 PDF-decode memory limit reached                                   0
 PDF-decode session limit reached                                  0
 PDF-decode malloc count                                           0
 PDF-decode free count                                             0
 PDF-decode bypassed  -  Global memory limit                       0
 PDF-decode bypassed  -  Per file memory limit                     0
 PDF-decode bypassed  -  Document encrypted                        0
 PDF-decode bypassed  -  Unsupported filter decode                 0
 PDF-decode bypassed  -  corrupted PDF file                        0
 PDF-decode bypassed  -  unsupported obj-stream length             0
 PDF-decode bypassed  -  filter decode failed                      0
```

show security idp counters pdf-decoder logical-system LSYS1

```
user@host>  show security idp counters pdf-decoder logical-system LSYS1
IDP counters:
```

```
 IDP counter type                                                Value
PDF-Decode enabled sessions                                       0
PDF-Decode requests                                               0
PDF-decode in pending state                                       0
PDF-decode finished sucessfully                                   0
PDF-decode per file memory limit reached                          0
PDF-decode memory limit reached                                   0
PDF-decode session limit reached                                  0
PDF-decode malloc count                                           0
PDF-decode free count                                             0
PDF-decode bypassed  -  Global memory limit                       0
PDF-decode bypassed  -  Per file memory limit                     0
PDF-decode bypassed  -  Document encrypted                        0
PDF-decode bypassed  -  Unsupported filter decode                 0
PDF-decode bypassed  -  corrupted PDF file                        0
PDF-decode bypassed  -  unsupported obj-stream length             0
PDF-decode bypassed  -  filter decode failed                      0
```

show security idp counters pdf-decoder tenant TSYS1

```
user@host>  show security idp counters pdf-decoder tenant TSYS1
IDP counters:

 IDP counter type                                                Value
PDF-Decode enabled sessions                                       0
PDF-Decode requests                                               0
PDF-decode in pending state                                       0
PDF-decode finished sucessfully                                   0
PDF-decode per file memory limit reached                          0
PDF-decode memory limit reached                                   0
PDF-decode session limit reached                                  0
PDF-decode malloc count                                           0
PDF-decode free count                                             0
PDF-decode bypassed  -  Global memory limit                       0
PDF-decode bypassed  -  Per file memory limit                     0
PDF-decode bypassed  -  Document encrypted                        0
PDF-decode bypassed  -  Unsupported filter decode                 0
PDF-decode bypassed  -  corrupted PDF file                        0
PDF-decode bypassed  -  unsupported obj-stream length             0
PDF-decode bypassed  -  filter decode failed                      0
```

Release Information

Command introduced for user logical systems in Junos OS Release 18.3R1

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

clear security idp counters pdf-decoder | 915

show security idp counters log

IN THIS SECTION

- Syntax | 1065
- Description | 1066
- Options | 1066
- Required Privilege Level | 1066
- Output Fields | 1066
- Sample Output | 1069
- Release Information | 1072

Syntax

```
show security idp counters log
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays the status of all IDP log counter values.

Options

none	Displays the status of all IDP log counter values.
logical-system *logical-system-name*	(Optional) Displays the status of all IDP log counter values for a specific logical system.
logical-system all	(Optional) Displays the status of all IDP log counter values for all logical systems.
tenant *tenant-name*	(Optional) Displays the status of all IDP log counter values for a specific tenant system.

Required Privilege Level

view

Output Fields

Table 59 on page 1066 lists the output fields for the `show security idp counters log` command. Output fields are listed in the approximate order in which they appear.

Table 59: show security idp counters log Output Fields

Field Name	Field Description
`Logs dropped`	Number of logs that are dropped.
`Suppressed log count`	Number of logs that are suppressed.

Table 59: show security idp counters log Output Fields *(Continued)*

Field Name	Field Description
`Logs waiting for post-window packets` (Unsupported)	Number of logs waiting for post-window packets.
`Logs ready to be sent` (Unsupported)	Number of logs ready to be sent.
`Logs in suppression list` (Unsupported)	Number of logs considered for suppression list.
`Log timers created`	Number of times the log timer is created.
`Logs timers expired`	Number of times the log timer is expired.
`Log timers cancelled`	Number of times the log timer is canceled.
`Logs ready to be sent high watermark` (Unsupported)	Number of packets that are ready to be sent with high degree watermark.
`Log receive buffer full` (Unsupported)	Number of times the buffer is full.
`Packet log too big` (Unsupported)	Number of packet logs that exceeded allowed packet log size.
`Reads per second` (Unsupported)	Number of packets that are read per second.

Table 59: show security idp counters log Output Fields *(Continued)*

Field Name	Field Description
`Logs in read buffer high watermark` (Unsupported)	Number of high watermark packets that are in read buffer.
`Packets logged`	Number of packets that are logged,
`Packets lost` (Unsupported)	Number of packets that are failed to log.
`Packets copied` (Unsupported)	Number of packets copied during packet log.
`Packets held` (Unsupported)	Number of packets held for packet log.
`Packets released`	Number of packets that are released from hold.
`IP Action Messages` (Unsupported)	Number of IP action messages.
`IP Action Drops` (Unsupported)	Number of IP action messages dropped.
`IP Action Exists` (Unsupported)	Number of exits during IP action creation.
`NWaits` (Unsupported)	Number of logs waiting for post window packets.

Table 59: show security idp counters log Output Fields *(Continued)*

Field Name	Field Description
Match vectors	Number of attacks in IDS match vector.
Supercedes	Number of attacks in supercede vector.

Sample Output

show security idp counters log

```
user@host> show security idp counters log
IDP counters:
IDP counter type                                                Value
Logs dropped                                                      0
Suppressed log count                                              0
Logs waiting for post-window packets                              0
Logs ready to be sent                                             0
Logs in suppression list                                          0
Log timers created                                                0
Logs timers expired                                               0
Log timers cancelled                                              0
Logs ready to be sent high watermark                              0
Log receive buffer full                                           0
Packet log too big                                                0
Reads per second                                                  1
Logs in read buffer high watermark                                0
Log Bytes in read buffer high watermark                           0
Packets logged                                                    0
Packets lost                                                      0
Packets copied                                                    0
Packets held                                                      0
Packets released                                                  0
IP Action Messages                                                0
IP Action Drops                                                   0
IP Action Exists                                                  0
```

```
NWaits                                                          0
Match vectors                                                   0
Supercedes                                                      0
Kpacket too big                                                 0
```

show security idp counters log logical-system LSYS1

```
user@host> show security idp counters log logical-system LSYS1
IDP counters:
IDP counter type                                              Value
Logs dropped                                                    0
Suppressed log count                                            0
Logs waiting for post-window packets                            0
Logs ready to be sent                                           0
Logs in suppression list                                        0
Log timers created                                              0
Logs timers expired                                             0
Log timers cancelled                                            0
Logs ready to be sent high watermark                            0
Log receive buffer full                                         0
Packet log too big                                              0
Reads per second                                                0
Logs in read buffer high watermark                              0
Log Bytes in read buffer high watermark                         0
Packets logged                                                  0
Packets lost                                                    0
Packets copied                                                  0
Packets held                                                    0
Packets released                                                0
IP Action Messages                                              0
IP Action Drops                                                 0
IP Action Exists                                                0
NWaits                                                          0
Match vectors                                                   0
Supercedes                                                      0
send succeed                                                    0
send fail                                                       0
retries on send failures                                        0
uac send succeed                                                0
uac send fail                                                   0
idpd to flowd alloc msg fail                                    0
```

```
idpd to flowd enqueue log  msg fail                                  0
idpd to flowd enqueue log msg succeed                                0
idpd to flowdlog msg dequeued                                        0
idpd to flowdlog unknown msg type                                    0
flowd send succeed                                                   0
flowd send fail                                                      0
objcache alloc failure for sc_pcap_mbuf_info_t                      0
pcap mbuf alloc fail counter                                         0
pcap mbuf reinj failed                                               0
pcap fragmented packets count                                        0
idpd to flowd pcap messages count in dedicated mode                  0
idpd pcap type1 messages count                                       0
idpd pcap type2 messages count                                       0
idpd pcap type3 messages count                                       0
Kpacket too big                                                      0
```

show security idp counters log tenant TSYS1

```
user@host> show security idp counters log tenant TSYS1
IDP counters:

  IDP counter type                                                  Value
 Logs dropped                                                        0
 Suppressed log count                                                0
 Logs waiting for post-window packets                                0
 Logs ready to be sent                                               0
 Logs in suppression list                                            0
 Log timers created                                                  0
 Logs timers expired                                                 0
 Log timers cancelled                                                0
 Logs ready to be sent high watermark                                0
 Log receive buffer full                                             0
 Packet log too big                                                  0
 Reads per second                                                    0
 Logs in read buffer high watermark                                  0
 Log Bytes in read buffer high watermark                             0
 Packets logged                                                      0
 Packets lost                                                        0
 Packets copied                                                      0
 Packets held                                                        0
 Packets released                                                    0
```

```
IP Action Messages                                              0
IP Action Drops                                                 0
IP Action Exists                                                0
NWaits                                                          0
Match vectors                                                   0
Supercedes                                                      0
send succeed                                                    1
send fail                                                       0
retries on send failures                                        0
uac send succeed                                                0
uac send fail                                                   0
idpd to flowd alloc msg fail                                    0
idpd to flowd enqueue log  msg fail                             0
idpd to flowd enqueue log msg succeed                           0
idpd to flowdlog msg dequeued                                   0
idpd to flowdlog unknown msg type                               0
flowd send succeed                                              0
flowd send fail                                                 0
 objcache alloc failure for sc_pcap_mbuf_info_t                 0
pcap mbuf alloc fail counter                                    0
pcap mbuf reinj failed                                          0
pcap fragmented packets count                                   0
idpd to flowd pcap messages count in dedicated mode             0
idpd pcap type1 messages count                                  0
idpd pcap type2 messages count                                  0
idpd pcap type3 messages count                                  0
Kpacket too big                                                 0
```

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

event-rate

clear security idp counters log

show security idp counters ips

IN THIS SECTION

- Syntax | 1073
- Description | 1073
- Options | 1073
- Required Privilege Level | 1074
- Output Fields | 1074
- Sample Output | 1077
- Release Information | 1080

Syntax

```
show security idp counters ips
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays the status of all IPS counter values.

Options

none	Displays the status of all IPS counter values for root-system.

logical-system *logical-system-name*	(Optional) Displays the status of all IPS counter values for a specific logical system.
logical-system all	(Optional) Displays the status of all IPS counter values for all logical systems.
tenant *tenant-name*	(Optional) Displays the status of all IPS counter values for a specific tenant system.

Required Privilege Level

view

Output Fields

Table 60 on page 1074 lists the output fields for the `show security idp counters ips` command. Output fields are listed in the approximate order in which they appear.

Table 60: show security idp counters ips Output Fields

Field Name	Field Description
`TCP fast path`	Number of TCP packets skipped for IDS processing.
`Layer-4 anomalies`	Number of Layer-4 protocol error or anomaly.
`Anomaly hash misses`	Number of times look failed on anomaly hash.
`Line context matches`	Number of attempts to match line based attacks in traffic stream.
`Stream256 context matches`	Number of attempts to match stream based attacks in first 256 bytes of traffic stream.
`Stream context matches`	Number of attempts to match stream based attacks in traffic stream.

Table 60: show security idp counters ips Output Fields *(Continued)*

Field Name	Field Description
Packet context matches	Number of attempts to match packet based attacks in traffic packet.
Packet header matches	Number of attempts to match packet header based attacks in traffic packet.
Context matches	Number of attempts to match protocol context based attacks in traffic stream.
Regular expression matches	Number of attempts to match PCRE expressions in traffic stream.
Tail DFAs	Number of attempts to match an attack on tail DFA group matches.
Exempted attacks	Number of attacks exempted from match as per exempt rulebase.
Out of order chains	Number of times attack is excluded from match due to member attacks in an attack group did not complete chain.
Partial chain matches	Number of attacks in partial chain match with attack scope as transaction.
IDS device FIFO size	Number of IDS contexts in virtual IDS device.
IDS device FIFO overflows	Number of times an IDS context can not be written as the IDS device is full.
Brute force queue size	Number of entries in the brute force queue.
IDS cache hits (Unsupported)	Number of sessions those found attack instance in IDS cache.
IDS cache misses (Unsupported)	Number of sessions those did not find attack instance in IDS cache.

Table 60: show security idp counters ips Output Fields *(Continued)*

Field Name	Field Description
Shellcode detection invocations	Number of times shell code match is attempted.
Wrong offsets	Number of times attack's offset is not within the service offset range.
No peer MAC (Unsupported)	Number of times flow peer MAC address is not available.
Number of times custom feed updated	Number of times the custom feeds are updated.
Number of times custom feed update failed due to error	Number of times the custom feed updates failed due to an error.
Number of times custom feed update failed due to out of memory	Number of times custom feed updates failed due to memory capacity.
Number of times custom feed update failed due to feed not found	Number of times custom feed updates failed due to the feed not found.
Number of times custom feed update returned unexpected value	Number of times custom feed updates returned an unexpected value.

Sample Output

show security idp counters ips

```
user@host> show security idp counters ips
IDP counters:
 IDP counter type                                                Value
 TCP fast path                                                     15
 Layer-4 anomalies                                                 0
 Anomaly hash misses                                               3
 Line context matches                                              5
 Stream256 context matches                                         5
 Stream context matches                                            5
 Packet context matches                                            0
 Packet header matches                                             0
 Context matches                                                  12
 Regular expression matches                                        0
 Tail DFAs                                                         0
 Exempted attacks                                                  0
 Out of order chains                                               0
 Partial chain matches                                             0
 IDS device FIFO size                                              0
 IDS device FIFO overflows                                         0
 Brute force queue size                                            0
 IDS cache hits                                                    0
 IDS cache misses                                                  0
 Shellcode detection invocations                                   0
 Wrong offsets                                                     0
 No peer MAC                                                       0
 Content-decompression memory usage in KB                          0
 Content-decompression memory over limit                           0
 Content-decompression gunzip called                               0
 Content-decompression gunzip failed                               0
 Content-decompression others called                               0
 Content-decompression others failed                               0
 Content-decompression input bytes                                 0
 Content-decompression output bytes                                0
 Content-decompression ratio over limit                            0
 Content-decompression type mismatch                               0
 URL track session bypassed                                        0
 Exceeded max Tail DFA transition limit                            0
```

```
 Number of times HS stream close failed                                0
 Number of times HS stream open failed                                 0
 Number of times HS scan stream failed                                 0
 Number of times HS scan failed                                        0
 Number of times custom feed updated                                   0
 Number of times custom feed update failed due to error                0
 Number of times custom feed update failed due to out of memory        0
 Number of times custom feed update failed due to feed not found       0
 Number of times custom feed update returned unexpected value          0
```

show security idp counters ips logical-system LSYS1

```
user@host> show security idp counters ips logical-system LSYS1
IDP counters:

  IDP counter type                                                   Value
 TCP fast path                                                        40
 Layer-4 anomalies                                                    0
 Anomaly hash misses                                                  4
 Line context matches                                                 0
 Stream256 context matches                                            0
 Stream context matches                                               0
 Packet context matches                                               0
 Packet header matches                                                0
 Context matches                                                      4
 Context reset                                                        0
 Regular expression matches                                           0
 Tail DFAs                                                            0
 Exempted attacks                                                     0
 Out of order chains                                                  0
 Partial chain matches                                                0
 IDS device FIFO size                                                 0
 IDS device FIFO overflows                                            0
 Brute force queue size                                               2
 IDS cache hits                                                       0
 IDS cache misses                                                     0
 Shellcode detection invocations                                      0
 Wrong offsets                                                        0
 No peer MAC                                                          0
 Content-decompression memory usage in KB                             0
 Content-decompression memory over limit                              0
```

```
  Content-decompression gunzip called                                  0
  Content-decompression gunzip failed                                  0
  Content-decompression others called                                  0
  Content-decompression others failed                                  0
  Content-decompression input bytes                                    0
  Content-decompression output bytes                                   0
  Content-decompression ratio over limit                               0
  Content-decompression type mismatch                                  0
  URL track session bypassed                                           0
  Exceeded max Tail DFA transition limit                               0
  Number of times HS stream close failed                               0
  Number of times HS stream open failed                                0
  Number of times HS scan stream failed                                0
  Number of times HS scan failed                                       0
  Number of times custom feed updated                                  0
  Number of times custom feed update failed due to error               0
  Number of times custom feed update failed due to out of memory       0
  Number of times custom feed update failed due to feed not found      0
  Number of times custom feed update returned unexpected value         0
```

show security idp counters ips tenant TSYS1

```
user@host> show security idp counters ips tenant TSYS1
IDP counters:

   IDP counter type                                                  Value
  TCP fast path                                                       16
  Layer-4 anomalies                                                   0
  Anomaly hash misses                                                 1
  Line context matches                                                0
  Stream256 context matches                                           0
  Stream context matches                                              0
  Packet context matches                                              0
  Packet header matches                                               0
  Context matches                                                     1
  Context reset                                                       0
  Regular expression matches                                          0
  Tail DFAs                                                           0
  Exempted attacks                                                    0
  Out of order chains                                                 0
  Partial chain matches                                               0
```

```
IDS device FIFO size                                                0
IDS device FIFO overflows                                           0
Brute force queue size                                              0
IDS cache hits                                                      0
IDS cache misses                                                    0
Shellcode detection invocations                                     0
Wrong offsets                                                       0
No peer MAC                                                         0
Content-decompression memory usage in KB                            0
Content-decompression memory over limit                             0
Content-decompression gunzip called                                 0
Content-decompression gunzip failed                                 0
Content-decompression others called                                 0
Content-decompression others failed                                 0
Content-decompression input bytes                                   0
Content-decompression output bytes                                  0
Content-decompression ratio over limit                              0
Content-decompression type mismatch                                 0
URL track session bypassed                                          0
Exceeded max Tail DFA transition limit                              0
Number of times HS stream close failed                              0
Number of times HS stream open failed                               0
Number of times HS scan stream failed                               0
Number of times HS scan failed                                      0
Number of times custom feed updated                                 0
Number of times custom feed update failed due to error              0
Number of times custom feed update failed due to out of memory      0
Number of times custom feed update failed due to feed not found     0
Number of times custom feed update returned unexpected value        0
```

Release Information

Command modified in Junos OS Release 11.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

ips

clear security idp counters ips

show security idp counters dfa

IN THIS SECTION

- Syntax | **1081**
- Description | **1081**
- Options | **1082**
- Required Privilege Level | **1082**
- Output Fields | **1082**
- Sample Output | **1083**
- Release Information | **1084**

Syntax

```
show security idp counters dfa
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays the status of all DFA counter values.

Options

none	Displays the status of all DFA counter values.
logical-system *logical-system-name*	(Optional) Displays the status of all DFA counter values for a specific logical system.
logical-system all	(Optional) Displays the status of all DFA counter values for all logical systems.
tenant *tenant-name*	(Optional) Displays the status of all DFA counter values for a specific tenant system.

Required Privilege Level

view

Output Fields

Table 61 on page 1082 lists the output fields for the `show security idp counters dfa` command. Output fields are listed in the approximate order in which they appear.

Table 61: show security idp counters dfa Output Fields

Field Name	Field Description
`DFA Group Merged Usage`	Number of DFA groups merged.
`DFA Matches`	Number of DFA matches found.

Sample Output

show security idp counters dfa

```
user@host> show security idp counters dfa
IDP counters:

  IDP counter type                                                  Value
 DFA Group Merged Usage                                              0
 DFA Matches                                                         0
 DFA compressed                                                      0
 DFA group compressed                                                0
 DFA uncompressed                                                    0
 DFA group uncompressed                                              0
                                                                     1
```

show security idp counters dfa logical-system LSYS1

```
user@host> show security idp counters dfa logical-system LSYS1
IDP counters:

  IDP counter type                                                  Value
 DFA Group Merged Usage                                              0
 DFA Matches                                                         0
 DFA compressed                                                      0
 DFA group compressed                                                0
 DFA uncompressed                                                    0
 DFA group uncompressed                                              0
```

show security idp counters dfa tenant TSYS1

```
user@host> show security idp counters dfa tenant TSYS1
IDP counters:

  IDP counter type                                                  Value
 DFA Group Merged Usage                                              0
 DFA Matches                                                         1
 DFA compressed                                                      0
```

```
DFA group compressed                                    0
DFA uncompressed                                        0
DFA group uncompressed                                  0
```

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

clear security idp counters dfa

show security idp counters flow

IN THIS SECTION

- Syntax | **1085**
- Description | **1085**
- Options | **1085**
- Required Privilege Level | **1085**
- Output Fields | **1086**
- Sample Output | **1093**
- Release Information | **1101**

Syntax

```
show security idp counters flow
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays the status of all IDP flow counter values.

NOTE: On SRX Series devices with IDP enabled, if IDP attacks are configured for a single direction (server or client), a flow in the opposite direction does not need IDP processing. For TCP traffic, the TCP optimization feature ensures minimal processing for these flows without running into reassembly errors.

Options

none	Displays the status of all IDP flow counter values.
logical-system *logical-system-name*	(Optional) Displays the status of all IDP flow counter values for a specific logical system.
logical-system all	(Optional) Displays the status of all IDP flow counter values for all logical systems.
tenant *tenant-name*	(Optional) Displays the status of all IDP flow counter values for a specific tenant system.

Required Privilege Level

view

Output Fields

Table 62 on page 1086 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 62: 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.

Table 62: show security idp counters flow Output Fields *(Continued)*

Field Name	Description
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 times Sessions exceed high mark	Number of times sessions exceeded the high mark.

Table 62: show security idp counters flow Output Fields *(Continued)*

Field Name	Description
`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.
`SM Sessions encountered memory failures`	Number of SM sessions that encountered memory failures.
`SM Packets on sessions with memory failures`	Number of SM packets that encountered memory failures.
`Sessions constructed`	Number of sessions established.
`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.

Table 62: show security idp counters flow Output Fields *(Continued)*

Field Name	Description
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.
Server-to-client flows tcp optimized	Number of server-to-client flow TCP sessions that are optimized.
Client-to-server flows tcp optimized	Number of client-to-server flow TCP sessions that are optimized.
Both directions flows ignored	Number of server-to-client and client-to-server flow sessions that are ignored.
Fail-over sessions dropped	Number of failover sessions dropped.

Table 62: 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 plug-in buffers dropped.
IDP Stream stbuf reinjected	Number of IDP stream plug-in buffers injected.

Table 62: show security idp counters flow Output Fields *(Continued)*

Field Name	Description
`Busy packets from stream plugin`	Number of packets saved as one or more packets of this session from stream plug-in.
`Busy packets from packets plugin`	Number of saved packets for IDP stream plug-in 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.
`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.

Table 62: show security idp counters flow Output Fields *(Continued)*

Field Name	Description
`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: • 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.
`IDP sessions ignored for content decompression in intel inspect mode`	Number of IDP session ignored for content decompression in the IDP intelligent inspection mode.

Table 62: show security idp counters flow Output Fields *(Continued)*

Field Name	Description
IDP sessions ignored for bytes depth limit in intel inspect mode	Number of IDP session ignored for bytes depth in the IDP intelligent inspection mode.
IDP sessions ignored for protocol decoding in intel inspect mode	Number of IDP session ignored for protocol decoding in the IDP intelligent inspection mode.
IDP sessions detected CPU usage crossed intel inspect CPU threshold	Number of IDP session detected when the CPU usage crosses the CPU threshold of the IDP intelligent inspection.
IDP sessions detected mem drop below intel inspect low mem threshold	Number of IDP session detected when memory drops below the IDP intelligent inspect low memory threshold.

Sample Output

show security idp counters flow

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

  IDP counter type                                                Value
 Fast-path packets                                                40252
 Slow-path packets                                                127
 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                                                92
Policy cache misses                                              67
Policy cache entries                                             67
Maximum flow hash collisions                                     0
Flow hash collisions                                             0
Gates added                                                      0
Gate matches                                                     0
Sessions deleted                                                 127
Sessions aged-out                                                0
Sessions in-use while aged-out                                   0
TCP flows marked dead on RST/FIN                                 13
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
IDP session gate creation requests                               0
IDP session gate creation acknowledgements                       0
IDP session gate hits                                            0
IDP session gate timeouts                                        0
Number of times Sessions crossed the CPU threshold value that is set    0
Number of times Sessions crossed the CPU upper threshold         0
Sessions constructed                                             127
SM Sessions ignored                                              0
SM Sessions dropped                                              0
SM Sessions interested                                           168
SM Sessions not interested                                       4
SM Sessions interest error                                       0
Sessions destructed                                              127
SM Session Create                                                127
SM Packet Process                                                52257
SM ftp data session ignored by idp                               0
SM Session close                                                 127
SM Client-to-server packets                                      20066
SM Server-to-client packets                                      32191
SM Client-to-server L7 bytes                                     167292
SM Server-to-client L7 bytes                                     28523514
Client-to-server flows ignored                                   1
Server-to-client flows ignored                                   1
Server-to-client flows tcp optimized                             3
Client-to-server flows tcp optimized                             0
```

```
Both directions flows ignored                                    32
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                                                   35155
kq op error                                                      0
kq op stop                                                       0
PME wait not set                                                 0
PME wait set                                                     0
PME KQ run not called                                            0
IDP sessions ignored for content decompression in intel inspect mode    0
IDP sessions ignored for bytes depth limit in intel inspect mode        0
IDP sessions ignored for protocol decoding in intel inspect mode        0
```

```
 IDP sessions detected CPU usage crossed intel inspect CPU threshold     0
 IDP sessions detected mem drop below intel inspect low mem threshold    0
```

show security idp counters flow tenant TSYS1

```
user@host> show security idp counters flow tenant TSYS1
IDP counters:

  IDP counter type                                                    Value
 Fast-path packets                                                     38
 Slow-path packets                                                     1
 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                                                   1
 Policy cache entries                                                  0
 Maximum flow hash collisions                                          0
 Flow hash collisions                                                  0
 Gates added                                                           0
 Gate matches                                                          0
 Sessions deleted                                                      1
 Sessions aged-out                                                     0
 Sessions in-use while aged-out                                        0
 TCP flows marked dead on RST/FIN                                      1
 Policy init failed                                                    0
 Policy reinit 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
 IDP session gate creation requests                                    0
 IDP session gate creation acknowledgements                            0
 IDP session gate hits                                                 0
 IDP session gate timeouts                                             0
```

```
Number of times Sessions crossed the CPU threshold value that is set    0
Number of times Sessions crossed the CPU upper threshold                0
Sessions constructed                                                    1
SM Sessions ignored                                                     0
SM Sessions dropped                                                     0
SM Sessions interested                                                  2
SM Sessions not interested                                              0
SM Sessions interest error                                              0
Sessions destructed                                                     1
SM Session Create                                                       1
SM Packet Process                                                       38
SM ftp data session ignored by idp                                      1
SM Session close                                                        1
SM Client-to-server packets                                             15
SM Server-to-client packets                                             23
SM Client-to-server L7 bytes                                            99
SM Server-to-client L7 bytes                                            367
Client-to-server flows ignored                                          0
Server-to-client flows ignored                                          0
Server-to-client flows tcp optimized                                    0
Client-to-server flows tcp optimized                                    0
Both directions flows ignored                                           1
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
NGAppID Events with no L7 App                                           0
NGAppID Events with no active-policy                                    0
NGAppID Detector failed from event handler                              0
NGAppID Detector failed from API                                        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                                                       37
 kq op error                                                           0
 kq op stop                                                            0
 PME wait not set                                                      0
 PME wait set                                                          0
 PME KQ run not called                                                 0
 IDP sessions ignored for content decompression in intel inspect mode  0
 IDP sessions ignored for bytes depth limit in intel inspect mode      0
 IDP sessions ignored for protocol decoding in intel inspect mode      0
 IDP sessions detected CPU usage crossed intel inspect CPU threshold   0
 IDP sessions detected mem drop below intel inspect low mem threshold  0
```

show security idp counters flow (unified parser)

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

  IDP counter type                                                    Value
 Fast-path packets                                                     38
 Slow-path packets                                                     1
 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                                                1
  Policy cache entries                                               0
  Maximum flow hash collisions                                       0
  Flow hash collisions                                               0
  Gates added                                                        0
  Gate matches                                                       0
  Sessions deleted                                                   1
  Sessions aged-out                                                  0
  Sessions in-use while aged-out                                     0
  TCP flows marked dead on RST/FIN                                   1
  Policy init failed                                                 0
  Policy reinit 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
  IDP session gate creation requests                                 0
  IDP session gate creation acknowledgements                         0
  IDP session gate hits                                              0
  IDP session gate timeouts                                          0
  Number of times Sessions crossed the CPU threshold value that is set    0
  Number of times Sessions crossed the CPU upper threshold           0
  Sessions constructed                                               1
  SM Sessions ignored                                                0
  SM Sessions dropped                                                0
  SM Sessions interested                                             2
  SM Sessions not interested                                         0
  SM Sessions interest error                                         0
  Sessions destructed                                               13
Sessions processed with Unified Parser                              11
  Sessions processed without Unified Parser                          4
  SM Session Create                                                 15
  SM Packet Process                                                 38
  SM ftp data session ignored by idp                                 1
  SM Session close                                                   1
  SM Client-to-server packets                                       15
  SM Server-to-client packets                                       23
  SM Client-to-server L7 bytes                                      99
  SM Server-to-client L7 bytes                                     367
  Client-to-server flows ignored                                     0
```

```
Server-to-client flows ignored                                  0
Server-to-client flows tcp optimized                            0
Client-to-server flows tcp optimized                            0
Both directions flows ignored                                   1
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
NGAppID Events with no L7 App                                   0
NGAppID Events with no active-policy                            0
NGAppID Detector failed from event handler                      0
NGAppID Detector failed from API                                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                                                  37
kq op error                                                     0
kq op stop                                                      0
PME wait not set                                                0
```

```
PME wait set                                                    0
PME KQ run not called                                           0
IDP sessions ignored for content decompression in intel inspect mode    0
IDP sessions ignored for bytes depth limit in intel inspect mode        0
IDP sessions ignored for protocol decoding in intel inspect mode        0
IDP sessions detected CPU usage crossed intel inspect CPU threshold     0
IDP sessions detected mem drop below intel inspect low mem threshold    0
```

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

flow (Security IDP)

clear security idp counters flow

show security idp counters http-decoder

IN THIS SECTION

- Syntax | **1102**
- Description | **1102**
- Options | **1102**
- Required Privilege Level | **1102**
- Output Fields | **1102**
- Sample Output | **1103**
- Release Information | **1104**

Syntax

```
show security idp counters http-decoder
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays the status of all HTTP decoders.

Options

none	Displays the status of all HTTP decoders.
logical-system *logical-system-name*	(Optional) Displays the status of all HTTP decoders for a specific logical system.
logical-system all	(Optional) Displays the status of all HTTP decoders for all logical systems.
tenant *tenant-name*	(Optional) Displays the status of all HTTP decoders for a specific tenant system.

Required Privilege Level

view

Output Fields

Table 63 on page 1103 lists the output fields for the `show security idp counters http-decoder` command. Output fields are listed in the approximate order in which they appear.

Table 63: show security idp counters http-decoder Output Fields

Field Name	Field Description
No of file-decoder requests from MIME over HTTP	Number of active file decoder requests sent over HTTP from MIME.
No of pending file-decoder requests from MIME over HTTP	Number of pending file decoder requests sent over HTTP from MIME.
No of completed file-decoder requests from MIME over HTTP	Number of completed file decoder requests sent over HTTP from MIME.
No of unrecognized file type from MIME over HTTP	Number of unrecognized file types sent over HTTP from MIME.
No of compressed payload transferred over HTTP	Number of compressed files transferred over HTTP from MIME.

Sample Output

show security idp counters http-decoder

```
user@host> show security idp counters http-decoder
IDP counters:
IDP counter type                                                Value
No of file-decoder requests from MIME over HTTP                   0
No of pending file-decoder requests from MIME over HTTP           0
No of completed file-decoder requests from MIME over HTTP         0
No of unrecognized file type from MIME over HTTP                  0
No of compressed payload transferred over HTTP                    0
No of bypassed files over HTTP                                  0
```

show security idp counters http-decoder logical-system LSYS1

```
user@host> show security idp counters http-decoder logical-system LSYS1
IDP counters:

  IDP counter type                                                  Value
 No of file-decoder requests from MIME over HTTP                     0
 No of pending file-decoder requests from MIME over HTTP             0
 No of completd file-decoder requests from MIME over HTTP            0
 No of unrecognized file type from MIME over HTTP                    0
 No of compressed payload transferred over HTTP                      0
 No of bypassed files over HTTP                                      0
```

show security idp counters http-decoder tenant TSYS1

```
user@host> show security idp counters http-decoder tenant TSYS1
IDP counters:

  IDP counter type                                                  Value
 No of file-decoder requests from MIME over HTTP                     0
 No of pending file-decoder requests from MIME over HTTP             0
 No of completd file-decoder requests from MIME over HTTP            0
 No of unrecognized file type from MIME over HTTP                    0
 No of compressed payload transferred over HTTP                      0
 No of bypassed files over HTTP                                      0
```

Release Information

Command introduced in Junos OS Release 11.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

clear security idp counters http-decoder

show security idp counters packet-log

IN THIS SECTION

- Syntax | 1105
- Description | 1105
- Options | 1106
- Required Privilege Level | 1106
- Output Fields | 1106
- Sample Output | 1110
- Release Information | 1119

Syntax

```
show security idp counters packet-log
<logical-system (logical-system-name | all)>
<tenant (tenant-name | all)>
```

Description

Displays the values of all IDP packet-log counters.

Starting in Junos OS Release 22.1R1, you can enable a secure SSL or TLS connection to send encrypted IDP packet capture log to the packet capture receiver. To establish the SSL or TLS connection, you must specify the SSL initiation profile that you want to use in the IDP packet log configuration. With this encryption support the output of this command is also enhanced to provide information on maximum number of RTCOM ports.

Options

none	Displays the values of all IDP packet-log counters.
logical-system *logical-system-name*	(Optional) Displays the values of all IDP packet-log counters for a specific logical system.
logical-system all	(Optional) Displays the values of all IDP packet-log counters for all logical systems. Displays values for all IDP counters from Junos release version 20.3R3.
tenant *tenant-name*	(Optional) Displays the values of all IDP packet-log counters for a specific tenant system.

Required Privilege Level

view

Output Fields

The following table lists the output fields for the `show security idp counters packet-log` command. Output fields are listed in the approximate order in which they appear.

Field Name	Field Description
`Total packets captured since packet capture was activated`	Number of packets captured by the device by the IDP service.
`Total sessions enabled since packet capture was activated`	Number of sessions that have performed packet capture since the capture facility was activated.
`Sessions currently enabled for packet capture`	Number of sessions that are actively capturing packets at this time.

(Continued)

Field Name	Field Description
Packets currently captured for enabled sessions	Number of packets that have been captured by active sessions.
Packet clone failures	Number of packet capture failures due to cloning error.
Session log object failures	Number of objects containing log messages generated during packet capture that were not successfully transmitted to the host.
Session packet log object failures	Number of objects containing captured packets that were not successfully transmitted to the host.
Sessions skipped because session limit exceeded	Number of sessions that could not initiate packet capture because the maximum number of sessions specified for the device were conducting captures at that time.
Packets skipped because packet limit exceeded	Number of packets not captured because the packet limit specified for this device was reached.
Packets skipped because total memory limit exceeded	Number of packets not captured because the memory allocated for packet capture on this device was exceeded.
Packet log host route lookup failures	Number of times captured packet failures due to missing route details.
Number of Packet-Log RTCOM create req successful	This counter is incremented each time RTCOM com create request is successful. This does not mean that the TCP + SSL handshake is complete, and connection is UP.
Number of Packet-Log RTCOM create req failed due to invalid param	This counter is incremented each time RTCOM com create request is failed due to an invalid parameter.

(Continued)

Field Name	Field Description
Number of Packet-Log RTCOM create req failed due to quota exceed	This counter is incremented each time RTCOM com create request is failed due to lack of ports availability.
Number of Packet-Log RTCOM create req failed due to status not ready	This counter is incremented each time RTCOM com create request is failed due to IDP plugin is not registered with RTCOM or registration was failed.
Number of Packet-Log RTCOM create done event received	This counter is incremented each time IDP plugin receives an event CREATE DONE from RTCOM plugin. This event is basically received after the handshake was complete and connection is established to the host.
Number of Packet-Log RTCOM peer close event received	This counter is incremented each time IDP plugin receives an event PEER CLOSE from RTCOM plugin. This event is basically received if the SSL/TLS connection is closed by host due to some reason.
Number of Packet-Log RTCOM conn abort event received	This counter is incremented each time IDP plugin receives an event CONN ABORT from RTCOM plugin. This event is basically received if the SSL/TLS connection is aborted due to some reason.
Number of Packet-Log RTCOM data rcvd event received	This counter is incremented each time IDP plugin receives an event DATA RCVD from RTCOM plugin. This event is basically received if some data is received on the established connection from the host.
Number of Packet-Log RTCOM connections established successfully	This counter is incremented when the SSL/TLS connection is successful to the host. Basically, handshake is completed, and data can be send over the established connection.
Number of Packet-Log RTCOM connections failed after retries	This counter is incremented when SSL/TLS connection can not be established to the specified host even after 3 retries. Only 3 connection attempts are made to the configured host initially and on receiving CONN ABORT or PEER CLOSE event. If connection is no successful even after 3 attempts, no more further attempts are performed and this counter is incremented.

(Continued)

Field Name	Field Description
Number of Packet-Log RTCOM connections closed successfully	Number of SSL/TLS connections closed successfully. This counter is incremented only for such connections on which connection is up and gets closed later.
Number of Packet-Log RTCOM data send successful	This counter is incremented each time RTCOM send is successful. For a single packet log, it will be incremented multiple time as packet log is send as a chunk of 1024 bytes.
Number of Packet-Log RTCOM data send failed	This counter is incremented each time RTCOM send is failed.
Number of Packet-Log RTCOM contexts created	This counter is incremented each time a new RTCOM context structure is created for a specific host ip, host port, source ip, LSYS identifier, SSL identifier combination.
Number of Packet-Log RTCOM contexts reused	This counter is incremented each time a RTCOM context structure is reused for a specific host IP, host port, source IP, LSYS identifier, SSL identifier combination. This happens when attacks are getting detected simultaneously on a specific LSYS across different sessions. All the packet logs corresponding to those different sessions are sent using same SSL/TLS connection.
Number of Packet-Log RTCOM contexts destroyed	This counter is incremented each time a RTCOM context structure is destroyed. This happens only when no session or no packet log is referring that context to send packet logs.
Number of Packet-Log RTCOM packet log send successful	This counter is incremented each time a packet log is send successfully.
Number of Packet-Log RTCOM packet log send failed	This counter is incremented each time a packet log is send failed.

Sample Output

show security idp counters packet-log

```
user@host> show security idp counters packet-log
IDP counters:                                                                 Value
 Total packets captured since packet capture was activated                     0
 Total sessions enabled since packet capture was activated                     0
 Sessions currently enabled for packet capture                                 0
 Packets currently captured for enabled sessions                               0
 Packet clone failures                                                         0
 Session log object failures                                                   0
 Session packet log object failures                                            0
 Sessions skipped because session limit exceeded                               0
 Packets skipped because packet limit exceeded                                 0
 Packets skipped because total memory limit exceeded                           0
 Packet log host route lookup failures                                         0
Number of Packet-Log RTCOM create req successful                              1
Number of Packet-Log RTCOM create req failed due to invalid param             0
Number of Packet-Log RTCOM create req failed due to quota exceed              0
Number of Packet-Log RTCOM create req failed due to status not ready          0
Number of Packet-Log RTCOM create done event received                         1
Number of Packet-Log RTCOM peer close event received                          0
Number of Packet-Log RTCOM conn abort event received                          0
Number of Packet-Log RTCOM data rcvd event received                           0
Number of Packet-Log RTCOM connections established successfully              1
Number of Packet-Log RTCOM connections failed after retries                   0
Number of Packet-Log RTCOM connections closed successfully                    1
Number of Packet-Log RTCOM data send successful                               20
Number of Packet-Log RTCOM data send failed                                   0
Number of Packet-Log RTCOM contexts created                                   1
Number of Packet-Log RTCOM contexts reused                                    9
Number of Packet-Log RTCOM contexts destroyed                                 1
Number of Packet-Log RTCOM packet log send successful                         10
Number of Packet-Log RTCOM packet log send failed                             0
```

show security idp counters packet-log logical-system LSYS1

```
user@host> show security idp counters packet-log logical-system LSYS1
IDP counters:
```

```
  IDP counter type                                                  Value
 No of file-decoder requests from MIME over HTTP                    0
 No of pending file-decoder requests from MIME over HTTP            0
 No of completd file-decoder requests from MIME over HTTP           0
 No of unrecognized file type from MIME over HTTP                   0
 No of compressed payload transferred over HTTP                     0
 No of bypassed files over HTTP                                     0
 Packet log host route lookup failures                              0
 Number of Packet-Log RTCOM create req successful                   1
 Number of Packet-Log RTCOM create req failed due to invalid param  0
 Number of Packet-Log RTCOM create req failed due to quota exceed   0
 Number of Packet-Log RTCOM create req failed due to status not ready 0
 Number of Packet-Log RTCOM create done event received              1
 Number of Packet-Log RTCOM peer close event received               0
 Number of Packet-Log RTCOM conn abort event received               0
 Number of Packet-Log RTCOM data rcvd event received                0
 Number of Packet-Log RTCOM connections established successfully   1
 Number of Packet-Log RTCOM connections failed after retries        0
 Number of Packet-Log RTCOM connections closed successfully         1
 Number of Packet-Log RTCOM data send successful                    20
 Number of Packet-Log RTCOM data send failed                        0
 Number of Packet-Log RTCOM contexts created                        1
 Number of Packet-Log RTCOM contexts reused                         9
 Number of Packet-Log RTCOM contexts destroyed                      1
 Number of Packet-Log RTCOM packet log send successful              10
 Number of Packet-Log RTCOM packet log send failed                  0

user@host> show security idp counters packet-log logical-system Tn1
IDP counters:

  IDP counter type                                                  Value
 Total packets captured since packet capture was activated          0
 Total sessions enabled since packet capture was activated          0
 Sessions currently enabled for packet capture                      0
 Packets currently captured for enabled sessions                    0
 Packet clone failures                                              0
 Session log object failures                                        0
 Session packet log object failures                                 0
 Sessions skipped because session limit exceeded                    0
 Packets skipped because packet limit exceeded                      0
 Packets skipped because total memory limit exceeded                0
Number of Packet-Log RTCOM create req successful                    1
```

```
 Number of Packet-Log RTCOM create req failed due to invalid param       0
 Number of Packet-Log RTCOM create req failed due to quota exceed        0
 Number of Packet-Log RTCOM create req failed due to status not ready    0
 Number of Packet-Log RTCOM create done event received                   1
 Number of Packet-Log RTCOM peer close event received                    0
 Number of Packet-Log RTCOM conn abort event received                    0
 Number of Packet-Log RTCOM data rcvd event received                     0
 Number of Packet-Log RTCOM connections established successfully        1
 Number of Packet-Log RTCOM connections failed after retries             0
 Number of Packet-Log RTCOM connections closed successfully              1
 Number of Packet-Log RTCOM data send successful                         20
 Number of Packet-Log RTCOM data send failed                             0
 Number of Packet-Log RTCOM contexts created                             1
 Number of Packet-Log RTCOM contexts reused                              9
 Number of Packet-Log RTCOM contexts destroyed                           1
 Number of Packet-Log RTCOM packet log send successful                   10
 Number of Packet-Log RTCOM packet log send failed                       0
```

show security idp counters packet-log tenant TSYS1

```
user@host> show security idp counters packet-log tenant TSYS1
IDP counters:

  IDP counter type                                                    Value
 Total packets captured since packet capture was activated             0
 Total sessions enabled since packet capture was activated             0
 Sessions currently enabled for packet capture                         0
 Packets currently captured for enabled sessions                       0
 Packet clone failures                                                 0
 Session log object failures                                           0
 Session packet log object failures                                    0
 Sessions skipped because session limit exceeded                       0
 Packets skipped because packet limit exceeded                         0
 Packets skipped because total memory limit exceeded                   0
Number of Packet-Log RTCOM create req successful                      1
 Number of Packet-Log RTCOM create req failed due to invalid param     0
 Number of Packet-Log RTCOM create req failed due to quota exceed      0
 Number of Packet-Log RTCOM create req failed due to status not ready  0
 Number of Packet-Log RTCOM create done event received                 1
 Number of Packet-Log RTCOM peer close event received                  0
 Number of Packet-Log RTCOM conn abort event received                  0
```

```
 Number of Packet-Log RTCOM data rcvd event received                    0
 Number of Packet-Log RTCOM connections established successfully       1
 Number of Packet-Log RTCOM connections failed after retries            0
 Number of Packet-Log RTCOM connections closed successfully             1
 Number of Packet-Log RTCOM data send successful                        20
 Number of Packet-Log RTCOM data send failed                            0
 Number of Packet-Log RTCOM contexts created                            1
 Number of Packet-Log RTCOM contexts reused                             9
 Number of Packet-Log RTCOM contexts destroyed                          1
 Number of Packet-Log RTCOM packet log send successful                  10
 Number of Packet-Log RTCOM packet log send failed                      0
```

show security idp counters packet-log tenant all

```
user@host> show security idp counters packet-log logical-system all
root@idpdevesx16-vsrx2-12> show security idp counters packet logical-system all



Logical System: root-logical-system
IDP counters:



IDP counter type Value
Processed packets 0
Dropped packets 0
Dropped ICMP packets 0
Dropped TCP packets 0
Dropped UDP packets 0
Dropped Other packets 0
Dropped by IDP Policy 0
Dropped by Error 0
Dropped sessions 0
Bad IP headers 0
Packets with IP options 0
Decapsulated packets 0
GRE decapsulations 0
PPP decapsulations 0
GTP decapsulations 0
```

```
GTP flows 0
TCP decompression uncompressed IP 0
TCP decompression compressed IP 0
Deferred-send packets 0
Ktimer entry optimized 0
TTL errors 0
Routing loops 0
STP drops 0
No-route packets 0
Flood IP 0
Invalid ethernet headers 0
Packets attached 0
IP Packet attach failed 0
Packets cloned 0
Packets allocated 0
Packets destructed 0
Packets destructed in pipeline 0
Packet data buffer allocated 0
Packet data buffer released 0
Buffer allocation on clone avoided 0
Late buffer allocation on clone 0
Distinct clone request 0
KPP clone buf cache allocated 0
KPP clone buf cache released 0
KPP clone buf cache used 0
KQMSG constructed 0
KQMSG destructed 0
KQMSG destructed in pipeline 0
jbuf copy failed 0
jbuf pullup failed 0
jbuf copy done 0
jbuf copy freed 0
jbuf copy reinjected 0



Logical System: LSYS0
IDP counters:



IDP counter type Value
Processed packets 0
```

```
Dropped packets 0
Dropped ICMP packets 0
Dropped TCP packets 0
Dropped UDP packets 0
Dropped Other packets 0
Dropped by IDP Policy 0
Dropped by Error 0
Dropped sessions 0
Bad IP headers 0
Packets with IP options 0
Decapsulated packets 0
GRE decapsulations 0
PPP decapsulations 0
GTP decapsulations 0
GTP flows 0
TCP decompression uncompressed IP 0
TCP decompression compressed IP 0
Deferred-send packets 0
Ktimer entry optimized 0
TTL errors 0
Routing loops 0
STP drops 0
No-route packets 0
Flood IP 0
Invalid ethernet headers 0
Packets attached 0
IP Packet attach failed 0
Packets cloned 0
Packets allocated 0
Packets destructed 0
Packets destructed in pipeline 0
Packet data buffer allocated 0
Packet data buffer released 0
Buffer allocation on clone avoided 0
Late buffer allocation on clone 0
Distinct clone request 0
KPP clone buf cache allocated 0
KPP clone buf cache released 0
KPP clone buf cache used 0
KQMSG constructed 0
KQMSG destructed 0
KQMSG destructed in pipeline 0
jbuf copy failed 0
```

```
jbuf pullup failed 0
jbuf copy done 0
jbuf copy freed 0
jbuf copy reinjected 0



Logical System: LSYS2
IDP counters:



IDP counter type Value
Processed packets 0
Dropped packets 0
Dropped ICMP packets 0
Dropped TCP packets 0
Dropped UDP packets 0
Dropped Other packets 0
Dropped by IDP Policy 0
Dropped by Error 0
Dropped sessions 0
Bad IP headers 0
Packets with IP options 0
Decapsulated packets 0
GRE decapsulations 0
PPP decapsulations 0
GTP decapsulations 0
GTP flows 0
TCP decompression uncompressed IP 0
TCP decompression compressed IP 0
Deferred-send packets 0
Ktimer entry optimized 0
TTL errors 0
Routing loops 0
STP drops 0
No-route packets 0
Flood IP 0
Invalid ethernet headers 0
Packets attached 0
IP Packet attach failed 0
Packets cloned 0
Packets allocated 0
```

```
Packets destructed 0
Packets destructed in pipeline 0
Packet data buffer allocated 0
Packet data buffer released 0
Buffer allocation on clone avoided 0
Late buffer allocation on clone 0
Distinct clone request 0
KPP clone buf cache allocated 0
KPP clone buf cache released 0
KPP clone buf cache used 0
KQMSG constructed 0
KQMSG destructed 0
KQMSG destructed in pipeline 0
jbuf copy failed 0
jbuf pullup failed 0
jbuf copy done 0
jbuf copy freed 0
jbuf copy reinjected 0



Logical System: LSYS1
IDP counters:

IDP counter type     Value
Processed packets     0
Dropped packets       0
Dropped ICMP packets  0
Dropped TCP packets   0
Dropped UDP packets   0
Dropped Other packets 0
Dropped by IDP Policy 0
Dropped by Error 0
Dropped sessions 0
Bad IP headers 0
Packets with IP options 0
Decapsulated packets 0
GRE decapsulations 0
PPP decapsulations 0
GTP decapsulations 0
GTP flows 0
TCP decompression uncompressed IP 0
TCP decompression compressed IP 0
```

```
Deferred-send packets 0
Ktimer entry optimized 0
TTL errors 0
Routing loops 0
STP drops 0
No-route packets 0
Flood IP 0
Invalid ethernet headers 0
Packets attached 0
IP Packet attach failed 0
Packets cloned 0
Packets allocated 0
Packets destructed 0
Packets destructed in pipeline 0
Packet data buffer allocated 0
Packet data buffer released 0
Buffer allocation on clone avoided 0
Late buffer allocation on clone 0
Distinct clone request 0
KPP clone buf cache allocated 0
KPP clone buf cache released 0
KPP clone buf cache used 0
KQMSG constructed 0
KQMSG destructed 0
KQMSG destructed in pipeline 0
jbuf copy failed 0
jbuf pullup failed 0
jbuf copy done 0
jbuf copy freed 0
jbuf copy reinjected 0

IDP counters:
LSYS1:
  IDP counter type                                                  Value
 No of file-decoder requests from MIME over HTTP                     0
 No of pending file-decoder requests from MIME over HTTP             0
 No of completd file-decoder requests from MIME over HTTP            0
 No of unrecognized file type from MIME over HTTP                    0
 No of compressed payload transferred over HTTP                      0
 No of bypassed files over HTTP                                      0
 Packet log host route lookup failures                               0
 Number of Packet-Log RTCOM create req successful                    1
 Number of Packet-Log RTCOM create req failed due to invalid param   0
```

```
Number of Packet-Log RTCOM create req failed due to quota exceed        0
Number of Packet-Log RTCOM create req failed due to status not ready    0
Number of Packet-Log RTCOM create done event received                   1
Number of Packet-Log RTCOM peer close event received                    0
Number of Packet-Log RTCOM conn abort event received                    0
Number of Packet-Log RTCOM data rcvd event received                     0
Number of Packet-Log RTCOM connections established successfully        1
Number of Packet-Log RTCOM connections failed after retries             0
Number of Packet-Log RTCOM connections closed successfully              1
Number of Packet-Log RTCOM data send successful                         20
Number of Packet-Log RTCOM data send failed                             0
Number of Packet-Log RTCOM contexts created                             1
Number of Packet-Log RTCOM contexts reused                              9
Number of Packet-Log RTCOM contexts destroyed                           1
Number of Packet-Log RTCOM packet log send successful                   10
Number of Packet-Log RTCOM packet log send failed                       0
```

Release Information

Command introduced in Junos OS Release 10.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

No Link Title

clear security idp counters packet-log | 934

show security idp counters packet

IN THIS SECTION

- Syntax | 1120
- Description | 1120
- Options | 1120
- Required Privilege Level | 1121
- Output Fields | 1121
- Sample Output | 1124
- Release Information | 1130

Syntax

```
show security idp counters packet
<logical-system (logical-system-name | all)>
<tenant (tenant-name | all)>
```

Description

Displays the status of all IDP packet counter values.

Options

none	Displays the status of all IDP packet counter values.
logical-system *logical-system-name*	(Optional) Displays the status of all IDP packet counter values for a specific logical system.

logical-system all	(Optional) Displays the status of all IDP packet counter values for all logical systems.
tenant *tenant-name*	(Optional) Displays the status of all IDP packet counter values for a specific tenant system.
tenant all	(Optional) Displays the status of all IDP packet counter values for a all tenant systems.

Required Privilege Level

view

Output Fields

Table 64 on page 1121 lists the output fields for the `show security idp counters packet` command. Output fields are listed in the approximate order in which they appear.

Table 64: show security idp counters packet Output Fields

Field Name	Field Description
`Processed packets`	Number of packets processed by the IDP service.
`Dropped packets`	Number of packets dropped by the IDP service. The counter for all dropped packets.
`Dropped by IDP policy`	Number of packets dropped by the IDP policy. The counter for dropped packets due to the action specified in the IDP policy (starting with the attack detection).

Table 64: show security idp counters packet Output Fields *(Continued)*

Field Name	Field Description
`Dropped by Error`	Number of packets dropped by error. The difference between `Dropped packets` and `Dropped by IDP policy`. IDS drops are primarily due to policy actions. Reassembly errors lead to packet drops. So all drops shown in `show security idp counters ips`, `show security idp counters flow` and `show security idp counters tcp-reassembler` add to `Dropped by Error`. All drops includes reassembly errors, anomalies similar to bad ip header and TTL errors.
`Dropped sessions` (Unsupported)	Number of sessions dropped.
`Bad IP headers`	Number of packets that fail IP header length validity check.
`Packets with IP options`	Number of packets that contain the optional header fields.
`Decapsulated packets`	Number of packets that are decapsulated.
`GRE decapsulations` (Unsupported)	Number of packets that are generic routing encapsulation (GRE) decapsulated.
`PPP decapsulations` (Unsupported)	Number of packets that are Point-to-Point Protocol (PPP) decapsulated.
`TCP decompression uncompressed IP` (Unsupported)	Number of uncompressed IP headers that are to be TCP decompressed.

Table 64: show security idp counters packet Output Fields *(Continued)*

Field Name	Field Description
`TCP decompression compressed IP` (Unsupported)	Number of compressed IP headers that are to be TCP decompressed.
`Deferred-send packets` (Unsupported)	Number of deferred IP packets that are sent out.
`IP-in-IP packets` (Unsupported)	Number of packets that are IP-in-IP encapsulated.
`TTL errors` (Unsupported)	Number of packets with TTL error in the header.
`Routing loops` (Unsupported)	Number of packets that continue to be routed in an endless circle due to an inconsistent routing state.
`No-route packets` (Unsupported)	Number of packets that could not be routed further.
`Flood IP` (Unsupported)	Number of packets that are identified as IP flood packets.
`Invalid ethernet headers` (Unsupported)	Number of packets that are identified with an invalid Ethernet header.
`Packets attached`	Number of packets attached.
`Packets cloned`	Number of packets that are cloned.

Table 64: show security idp counters packet Output Fields *(Continued)*

Field Name	Field Description
`Packets allocated`	Number of packets allocated.
`Packets destructed`	Number of packets destructed.

Sample Output

show security idp counters packet

```
user@host> show security idp counters packet
IDP counters:
 IDP counter type                                                Value
 Processed packets                                                 27
 Dropped packets                                                    0
 Dropped by IDP policy                                              0
 Dropped by error                                                   0
 Dropped sessions                                                   0
 Bad IP headers                                                     0
 Packets with IP options                                            0
 Decapsulated packets                                               0
 GRE decapsulations                                                 0
 PPP decapsulations                                                 0
 TCP decompression uncompressed IP                                  0
 TCP decompression compressed IP                                    0
 Deferred-send packets                                              0
 IP-in-IP packets                                                   0
 TTL errors                                                         0
 Routing loops                                                      0
 STP drops                                                          0
 No-route packets                                                   0
 Flood IP                                                           0
 Invalid ethernet headers                                           0
 Packets attached                                                  28
 Packets cloned                                                    28
```

```
 Packets allocated                                                0
 Packets destructed                                               55
```

show security idp counters packet logical-system LSYS1

```
user@host> show security idp counters packet logical-system LSYS1
IDP counters:
Logical System: LSYS1
  IDP counter type                                               Value
 Processed packets                                                64
 Dropped packets                                                  0
 Dropped ICMP packets                                             0
 Dropped TCP packets                                              0
 Dropped UDP packets                                              0
 Dropped Other packets                                            0
 Dropped by IDP Policy                                            0
 Dropped by Error                                                 0
 Dropped sessions                                                 0
 Bad IP headers                                                   0
 Packets with IP options                                          0
 Decapsulated packets                                             0
 GRE decapsulations                                               0
 PPP decapsulations                                               0
 GTP decapsulations                                               0
 GTP flows                                                        0
 TCP decompression uncompressed IP                                0
 TCP decompression compressed IP                                  0
 Deferred-send packets                                            0
 IP-in-IP packets                                                 0
 TTL errors                                                       0
 Routing loops                                                    0
 STP drops                                                        0
 No-route packets                                                 0
 Flood IP                                                         0
 Invalid ethernet headers                                         0
 Packets attached                                                64
 IP Packet attach failed                                          0
 Packets cloned                                                  25
 Packets allocated                                                0
 Packets destructed                                              89
 Packet data buffer allocated                                    24
```

```
  Packet data buffer released                                    24
  Buffer allocation on clone avoided                             0
  Late buffer allocation on clone                                0
  Distinct clone request                                         0
  KPP clone buf cache allocated                                  0
  KPP clone buf cache released                                   0
  KPP clone buf cache used                                       0
  KQMSG constructed                                              69
  KQMSG destructed                                               69
  jbuf copy failed                                               0
  jbuf pullup failed                                             0
  jbuf copy done                                                 0
  jbuf copy freed                                                0
  jbuf copy reinjected                                           0
```

show security idp counters packet tenant TSYS1

```
user@host> show security idp counters packet tenant TSYS1
IDP counters:
Tenant: TSYS1
   IDP counter type                                             Value
  Processed packets                                              38
  Dropped packets                                                0
  Dropped ICMP packets                                           0
  Dropped TCP packets                                            0
  Dropped UDP packets                                            0
  Dropped Other packets                                          0
  Dropped by IDP Policy                                          0
  Dropped by Error                                               0
  Dropped sessions                                               0
  Bad IP headers                                                 0
  Packets with IP options                                        0
  Decapsulated packets                                           0
  GRE decapsulations                                             0
  PPP decapsulations                                             0
  GTP decapsulations                                             0
  GTP flows                                                      0
  TCP decompression uncompressed IP                              0
  TCP decompression compressed IP                                0
  Deferred-send packets                                          0
  IP-in-IP packets                                               0
```

```
TTL errors                                                    0
Routing loops                                                 0
STP drops                                                     0
No-route packets                                              0
Flood IP                                                      0
Invalid ethernet headers                                      0
Packets attached                                              38
IP Packet attach failed                                       0
Packets cloned                                                21
Packets allocated                                             0
Packets destructed                                            59
Packets destructed in pipeline                                0
Packet data buffer allocated                                  21
Packet data buffer released                                   21
Buffer allocation on clone avoided                            0
Late buffer allocation on clone                               0
Distinct clone request                                        0
KPP clone buf cache allocated                                 0
KPP clone buf cache released                                  0
KPP clone buf cache used                                      0
KQMSG constructed                                             38
KQMSG destructed                                              38
KQMSG destructed in pipeline                                  0
jbuf copy failed                                              0
jbuf pullup failed                                            0
jbuf copy done                                                0
jbuf copy freed                                               0
jbuf copy reinjected                                          0
```

show security idp counters packet tenant all

```
Tenant: Tn3
IDP counters:

  IDP counter type                                            Value
Processed packets                                             0
Dropped packets                                               0
Dropped ICMP packets                                          0
Dropped TCP packets                                           0
Dropped UDP packets                                           0
Dropped Other packets                                         0
Dropped by IDP Policy                                         0
```

```
Dropped by Error                                                0
Dropped sessions                                                0
Bad IP headers                                                  0
Packets with IP options                                         0
Decapsulated packets                                            0
GRE decapsulations                                              0
PPP decapsulations                                              0
GTP decapsulations                                              0
GTP flows                                                       0
TCP decompression uncompressed IP                               0
TCP decompression compressed IP                                 0
Deferred-send packets                                           0
Ktimer entry optimized                                          0
TTL errors                                                      0
Routing loops                                                   0
STP drops                                                       0
No-route packets                                                0
Flood IP                                                        0
Invalid ethernet headers                                        0
Packets attached                                                0
IP Packet attach failed                                         0
Packets cloned                                                  0
Packets allocated                                               0
Packets destructed                                              0
Packets destructed in pipeline                                  0
Packet data buffer allocated                                    0
Packet data buffer released                                     0
Buffer allocation on clone avoided                              0
Late buffer allocation on clone                                 0
Distinct clone request                                          0
KPP clone buf cache allocated                                   0
KPP clone buf cache released                                    0
KPP clone buf cache used                                        0
KQMSG constructed                                               0
KQMSG destructed                                                0
KQMSG destructed in pipeline                                    0
jbuf copy failed                                                0
jbuf pullup failed                                              0
jbuf copy done                                                  0
jbuf copy freed                                                 0
jbuf copy reinjected                                            0

Tenant: Tn2
```

```
IDP counters:

  IDP counter type                                                Value
Processed packets                                                 0
Dropped packets                                                   0
Dropped ICMP packets                                              0
Dropped TCP packets                                               0
Dropped UDP packets                                               0
Dropped Other packets                                             0
Dropped by IDP Policy                                             0
Dropped by Error                                                  0
Dropped sessions                                                  0
Bad IP headers                                                    0
Packets with IP options                                           0
Decapsulated packets                                              0
GRE decapsulations                                                0
PPP decapsulations                                                0
GTP decapsulations                                                0
GTP flows                                                         0
TCP decompression uncompressed IP                                 0
TCP decompression compressed IP                                   0
Deferred-send packets                                             0
Ktimer entry optimized                                            0
TTL errors                                                        0
Routing loops                                                     0
STP drops                                                         0
No-route packets                                                  0
Flood IP                                                          0
Invalid ethernet headers                                          0
Packets attached                                                  0
IP Packet attach failed                                           0
Packets cloned                                                    0
Packets allocated                                                 0
Packets destructed                                                0
Packets destructed in pipeline                                    0
Packet data buffer allocated                                      0
Packet data buffer released                                       0
Buffer allocation on clone avoided                                0
Late buffer allocation on clone                                   0
Distinct clone request                                            0
KPP clone buf cache allocated                                     0
KPP clone buf cache released                                      0
KPP clone buf cache used                                          0
```

```
KQMSG constructed                                                0
KQMSG destructed                                                 0
KQMSG destructed in pipeline                                     0
jbuf copy failed                                                 0
jbuf pullup failed                                               0
jbuf copy done                                                   0
jbuf copy freed                                                  0
jbuf copy reinjected                                             0
```

Usage

The command "show security idp counters" with options such as action, application-identification, dfa, flow, http-decoder, ips, log, memory, packet, packet-log, pdf-decoder, policy-manager, tcp-reassembler, now displays system names for all IDP counters in addition to root-logical-system counters.

Release Information

Command introduced in Junos OS Release 9.2.

The fields `Dropped by IDP policy` and `Dropped by Error` added in Junos OS Release 10.1.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

clear security idp counters packet | 936

show security idp counters policy-manager

IN THIS SECTION

- Syntax | 1131
- Description | 1131
- Options | 1131

- Required Privilege Level | 1132
- Output Fields | 1132
- Sample Output | 1132
- Release Information | 1133

Syntax

```
show security idp counters policy-manager
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays the status of all IDP policies counter values.

Options

none	Displays the status of all IDP policies counter values.
logical-system *logical-system-name*	(Optional) Displays the status of all IDP policies counter values for a specific logical system.
logical-system all	(Optional) Displays the status of all IDP policies counter values for all logical systems.
tenant *tenant-name*	(Optional) Displays the status of all IDP policies counter values for a specific tenant system.

Required Privilege Level

view

Output Fields

Table 65 on page 1132 lists the output fields for the `show security idp counters policy-manager` command. Output fields are listed in the approximate order in which they appear.

Table 65: show security idp counters policy-manager Output Fields

Field Name	Field Description
`Number of policies`	Number of policies installed.
`Number of aged out policies`	Number of IDP policies that are expired.

Sample Output

show security idp counters policy-manager

```
user@host> show security idp counters policy-manager
IDP counters:
 IDP counter type                                                Value
 Number of policies                                                0
 Number of aged out policies                                       0
```

show security idp counters policy-manager logical-system LSYS1

```
user@host> show security idp counters policy-manager logical-system LSYS1
IDP counters:

  IDP counter type                                                Value
```

```
 Number of policies                                                  1
 Number of aged out policies                                         0
 Policy compile failure due to memory                                0
```

show security idp counters policy-manager tenant TSYS1

```
user@host> show security idp counters policy-manager tenant TSYS1
IDP counters:

  IDP counter type                                                 Value
 Number of policies                                                  0
 Number of aged out policies                                         0
 Policy compile failure due to memory                                0
```

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

clear security idp counters policy-manager

show security idp counters tcp-reassembler

IN THIS SECTION

- Syntax | **1134**
- Description | **1134**

- Options | 1134
- Required Privilege Level | 1135
- Output Fields | 1135
- Sample Output | 1138
- Release Information | 1140

Syntax

```
show security idp counters tcp-reassembler
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays the status of all TCP reassembler counter values.

NOTE: On SRX Series devices with IDP enabled, if IDP attacks are configured for a single direction (server or client), a flow in the opposite direction does not need IDP processing. For TCP traffic, the TCP optimization feature ensures minimal processing for these flows without running into reassembly errors.

Options

none	Displays the status of all TCP reassembler counter values.
logical-system *logical-system-name*	(Optional) Displays the status of all TCP reassembler counter values for a specific logical system.
logical-system all	(Optional) Displays the status of all TCP reassembler counter values for all logical systems.

tenant *tenant-name* (Optional) Displays the status of all TCP reassembler counter values for a specific tenant system.

Required Privilege Level

view

Output Fields

Table 66 on page 1135 lists the output fields for the `show security idp counters tcp-reassembler` command. Output fields are listed in the approximate order in which they appear.

Table 66: show security idp counters tcp-reassembler Output Fields

Field Name	Field Description
`Bad TCP checksums` (Unsupported)	Number of packets that have incorrect TCP checksums.
`Bad TCP headers`	Number of bad TCP headers detected.
`Slow path segments`	Number of segments that are sent through the slow path if the TCP segment does not pass fast-path segment validation.
`Fast path segments`	Number of segments that are sent through the fast path after passing a predefined TCP validation sequence.
`Tcp Optimized s2c segments`	Number of TCP segments that are sent through optimized re-assembly process from server to client.
`Tcp Optimized c2s segments`	Number of TCP segments that are sent through optimized re-assembly process from server to client.

Table 66: show security idp counters tcp-reassembler Output Fields *(Continued)*

Field Name	Field Description
`Sequence number wrap around errors`	Number of packets that wrap around of the sequence number.
`Session reuses`	Number of sessions that reused an already established TCP session.
`SYN retransmissions`	Number of SYN packets that are retransmitted.
`Bad three way handshake acknowledgements`	Number of packets that have incorrect three-way handshake acknowledgements (ACK packet).
`Sequence number out of sync flows`	Number of packets that have out-of-sync sequence numbers.
`Fast path pattern matches in queued up streams`	Number of queued packets that have fast path pattern match.
`New segments with no overlaps with old segment`	Number of new segments that do not overlap with old segment.
`New segment overlaps with beginning of old segment`	Number of new segments that overlap with beginning of old segment.
`New segment overlaps completely with old segment`	Number of new segments that overlap completely with old segment.
`New segment is contained in old segment`	Number of new segments contained in old segment.
`New segment overlaps with end of old segment`	Number of new segments that overlap with the end of old segment.

Table 66: show security idp counters tcp-reassembler Output Fields *(Continued)*

Field Name	Field Description
New segment begins after end of old segment	Number of new segments that overlap after the end of old segment.
Memory consumed by new segment	Memory that is consumed by the new segment.
Peak memory consumed by new segments	Peak memory that is consumed by the new segment.
Segments in memory	Number of segments that are stored in memory for processing.
Per-flow memory overflows	Number of segments dropped after reaching per flow memory limit.
Global memory overflows	Number of segments dropped after reaching reassembler global memory limit.
Overflow drops	Number of packets that are dropped due to memory overflow.
Copied packets (Unsupported)	Number of packets copied in reassembler.
Closed Acks	Number of Ack packets seen without having seen SYN on the same session.
Ack Validation failures	Number of Invalid ACKs received from server during 3-way handshake.
Simultanious syn	Number of simultaneous syn packets seen.
C2S synack	Number of C2S Syn/Ack packets seen.
Segment to left of receiver window	Number of segments falling left of receive window.

Table 66: show security idp counters tcp-reassembler Output Fields *(Continued)*

Field Name	Field Description
Segment to right of receiver window	Number of segments falling right of receive window.
SYN seen in the window	Number of Syn packets seen after connection establishment.
ACK bit is off	Number of packets seen without ACK after connection establishment.
Unexpected FIN	Number of unexpected FIN packets seen.
Duplicate Syn/Ack with different SEQ	Number of Syn/Ack packets with different SEQ numbers.

Sample Output

show security idp counters tcp-reassembler

```
user@host> show security idp counters tcp-reassembler
 IDP counters:

  IDP counter type                                                Value
 Bad TCP checksums                                                 0
 Bad TCP headers                                                   0
 Slow path segments                                                90
 Fast path segments                                                7099
 Tcp Optimized s2c segments                                        0
 Tcp Optimized c2s segments                                        0
 Sequence number wrap around errors                                0
 Session reuses                                                    0
 SYN retransmissions                                               0
 Bad three way handshake acknowledgements                          0
 Sequence number out of sync flows                                 0
 Fast path pattern matches in queued up streams                    0
```

```
  New segments with no overlaps with old segment                    0
  New segment overlaps with beginning of old segment                0
  New segment overlaps completely with old segment                  0
  New segment is contained in old segment                           0
  New segment overlaps with end of old segment                      0
  New segment begins after end of old segment                       3
  Memory consumed by new segment                                    0
  Peak memory consumed by new segments                              3821
  Segments in memory                                                0
  Per-flow memory overflows                                         0
  Global memory overflows                                           0
  Overflow drops                                                    0
  Copied packets                                                    0
  Closed Acks                                                       3
  Ack Validation failure                                            0
  Simultanious syn                                                  0
  C2S synack                                                        0
  segment to left of receiver window                                0
  segment to right of receiver window                               0
  SYN seen in the window                                            0
  ACK bit is off                                                    0
  Unexpected FIN                                                    0
  Duplicate Syn/Ack with different SEQ                              0
```

show security idp counters tcp-reassembler logical-system LSYS1

```
user@host> show security idp counters tcp-reassembler logical-system LSYS1
IDP counters:

   IDP counter type                                                 Value
  Bad TCP checksums                                                  0
  Bad TCP headers                                                    0
  Slow path segments                                                 37
  Fast path segments                                                 27
  Tcp Optimized s2c segments                                         0
  Tcp Optimized c2s segments                                         0
  Sequence number wrap around errors                                 0
  Session reuses                                                     0
  SYN retransmissions                                                0
  Bad three way handshake acknowledgements                           0
  Sequence number out of sync flows                                  0
```

```
Fast path pattern matches in queued up streams                  0
New segments with no overlaps with old segment                  0
New segment overlaps with beginning of old segment              0
New segment overlaps completely with old segment                0
New segment is contained in old segment                         0
New segment overlaps with end of old segment                    0
New segment begins after end of old segment                     0
Memory consumed by new segment                                  0
Peak memory consumed by new segments                            2021
Segments in memory                                              0
Per-flow memory overflows                                       0
Global memory overflows                                         0
Overflow drops                                                  0
Overflow drops - missing packets                                0
Copied packets                                                  0
Closed Acks                                                     0
Ack Validation failure                                          0
Simultanious syn                                                0
C2S synack                                                      0
segment to left of receiver window                              0
segment to right of receiver window                             0
SYN seen in the window                                          0
ACK bit is off                                                  0
Unexpected FIN                                                  0
Duplicate Syn/Ack with different SEQ                            0
```

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

re-assembler

clear security idp counters tcp-reassembler

show security idp logical-system policy-association

IN THIS SECTION

- Syntax | 1141
- Description | 1141
- Required Privilege Level | 1141
- Output Fields | 1141
- Sample Output | 1142
- Release Information | 1142

Syntax

```
show security idp logical-system policy-association
```

Description

Display the IDP policy assigned to a logical system. The IDP policy is assigned to a logical system through the security profile.

Required Privilege Level

view

Output Fields

Table 67 on page 1142 lists the output fields for the `show security idp logical-system policy-association` command.

Table 67: show security idp logical-system policy-association Output Fields

Field Name	Field Description
`Logical system`	Name of the logical system to which an IDP policy is assigned.
`IDP policy`	Name of the IDP policy that is specified in the security profile that is bound to the logical system.

Sample Output

show security idp logical-system policy-association

```
user@host> show security idp logical-system policy-association
Logical system         IDP policy
root-logical-system    idp-policy1
lsys1                  idp-policy2
```

Release Information

Command introduced in Junos OS Release 11.3.

RELATED DOCUMENTATION

security-profile | 871

show security idp policies

IN THIS SECTION

- Syntax | 1143
- Description | 1143
- Options | 1143
- Required Privilege Level | 1144
- Output Fields | 1144
- Sample Output | 1144
- Release Information | 1145

Syntax

```
show security idp policies
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays the list of currently installed policies.

Options

none	Displays the list of currently installed policies.
logical-system *logical-system-name*	(Optional) Displays the list of currently installed policies for a specific logical system.

logical-system all	(Optional) Displays the list of currently installed policies for all logical systems.
tenant *tenant-name*	(Optional) Displays the list of currently installed policies for a specific tenant system.

Required Privilege Level

view

Output Fields

Sample Output

show security idp policies

```
user@host>show security idp policies
PIC : FPC 0 PIC 0:
  ID    Name                   Sessions    Memory      Detector
      0        idp-policy-unified         0        10179      12.6.130180509
```

show security idp policies logical-system LSYS0

```
user@host> show security idp policies logical-system LSYS0
 PIC : FPC 0 PIC 0:
ID    Name                   Sessions    Memory      Detector
53      idp_one policy             0           189712      12.6.130180509
```

Release Information

Command introduced in Junos OS Release 10.1.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show security idp active-policy

show security idp policy-commit-status

IN THIS SECTION

- Syntax | **1145**
- Description | **1146**
- Options | **1146**
- Required Privilege Level | **1146**
- Output Fields | **1146**
- Sample Output | **1146**
- Sample Output | **1147**
- Sample Output | **1147**
- Release Information | **1147**

Syntax

```
show security idp policy-commit-status
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays the IDP policy commit status. For example, status of policy compilation or load.

Options

none	Displays the IDP policy commit status.
logical-system *logical-system-name*	(Optional) Displays the IDP policy commit status for a specific logical system.
logical-system all	(Optional) Displays the IDP policy commit status for all logical systems.
tenant *tenant-name*	(Optional) Displays the IDP policy commit status for a specific tenant system.

Required Privilege Level

view

Output Fields

Sample Output

show security idp policy-commit-status

```
user@host> show security idp policy-commit-status
IDP policy[/var/db/idpd/bins/test.bin.gz.v] and detector[/var/db/idpd/sec-repository/installed-
detector/libidp-detector.so.tgz.v] loaded successfully.

The loaded policy size is:45583070 Bytes
```

Sample Output

show security idp policy-commit-status (on vSRX when you configure dynamic attack groups filters)

```
user@host> show security idp policy-commit-status
 Last good policy file does not exist. Aborted
```

Sample Output

show security idp policy-commit-status logical-system LSYS1

```
user@host> show security idp policy-commit-status logical-system LSYS1
 IDP policy[/var/db/idpd/bins//idp-policy-combined.bin.gz.v] and detector[/var/db/idpd/sec-
repository/installed-detector/libidp-detector.so.tgz.v] loaded successfully.
 The loaded policy size is:7416 Bytes
```

Release Information

Command introduced in JUNOS OS Release 10.4.

Starting with Junos OS Release 12.3X48-D15 and Junos OS Release 17.3R1, a new pattern matching engine is introduced for the SRX Series IDP feature. This scanning mechanism helps improve performance and policy loading. The new engine is 9.223 times faster than the existing DFA engine.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

show security idp status

show security idp policy-commit-status clear

show security idp policy-templates-list

IN THIS SECTION

- Syntax | 1148
- Description | 1148
- Required Privilege Level | 1148
- Output Fields | 1148
- Sample Output | 1149
- Release Information | 1149

Syntax

```
show security idp policy-templates-list
```

Description

This command displays the list of available policy templates for logical systems.

Required Privilege Level

view

Output Fields

Sample Output

show security idp policy-templates-list

```
user@host>show security idp policy-templates-list
 Web_Server
DMZ_Services
DNS_Service
File_Server
Getting_Started
IDP_Default
Server-Protection
Server-Protection-1G
Client-Protection
Client-Protection-1G
Client-And-Server-Protection
Client-And-Server-Protection-1G
Recommended
```

Release Information

Command introduced in Junos OS Release 10.1.

Command introduced for user logical system in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

show security idp active-policy

show security idp security-package-version

IN THIS SECTION

- Syntax | 1150
- Description | 1150
- Options | 1150
- Required Privilege Level | 1151
- Output Fields | 1151
- Sample Output | 1152
- Release Information | 1152

Syntax

```
show security idp security-package-version
<logical-system (logical-system-name | all)>
<tenant tenant-name>
```

Description

Displays information of the currently installed security package version and detector version.

Options

none	Displays information of the currently installed security package version and detector version.
logical-system *logical-system-name*	(Optional) Displays information of the currently installed security package version and detector version for a specific logical system.

logical-system all	(Optional) Displays information of the currently installed security package version and detector version for all logical systems.
tenant *tenant-name*	(Optional) Displays information of the currently installed security package version and detector version for a specific tenant system.

Required Privilege Level

view

Output Fields

Table 68 on page 1151 lists the output fields for the `show security idp security-package-version` command. Output fields are listed in the approximate order in which they appear.

Table 68: show security idp security-package-version Output Fields

Field Name	Field Description
`Attack database version`	Attack database version number that is currently installed on the system.
`Detector version`	Detector version number that is currently installed on the system.
`Policy template version`	Policy template version number that is currently installed on the system.
`Rollback version`	(SRX) Rollback version number on the system.
`Rollback Detector version`	Rollback Detector version number on the system.

Sample Output

show security idp security-package-version (Output After Initial Bootup)

```
user@host> show security idp security-package-version
Attack database version:N/A(N/A)
Detector version :N/A
Policy template version :N/A
Rollback Attack database version :N/A(N/A)
Rollback Detector version : N/A
```

show security idp security-package-version (Output After Security Package is Installed)

```
user@host> show security idp security-package-version
 Attack database version :3525(Tue Sep 13 17:40:07 2022 UTC)
  Detector version :12.6.180220128
  Policy template version :N/A
  Rollback Attack database version :N/A(N/A)
  Rollback Detector version :N/A
```

show security idp security-package-version

```
user@host> show security idp security-package-version
  Attack database version :3533(Mon Oct 10 14:57:03 2022 UTC)
  Detector version :12.6.180220128
  Policy template version :3521
  Rollback Attack database version :3527(Tue Sep 20 14:21:36 2022 UTC)
  Rollback Detector version :12.6.180220128
```

Release Information

Command introduced in Junos OS Release 9.2.

`logical-system` option introduced in Junos OS Release 18.3R1.

`tenant` option introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

- No Link Title
- No Link Title
- No Link Title

show security ike security-associations

IN THIS SECTION

- Syntax | 1153
- Description | 1154
- Options | 1154
- Required Privilege Level | 1155
- Output Fields | 1155
- Sample Output | 1161
- Release Information | 1170

Syntax

```
show security ike security-associations
<peer-address>
<brief | detail>
<family (inet  | inet6)>
<fpc slot-number>
<index SA-index-number>
<kmd-instance (all | kmd-instance-name)>
<pic slot-number>
<sa-type shortcut >
```

```
<srg-id id-number>
<ha-link-encryption>
```

Description

Display information about Internet Key Exchange security associations (IKE SAs).

Options

- none—Display standard information about existing IKE SAs, including index numbers.
- *peer-address*—(Optional) Display details about a particular SA based on the IPv4 or IPv6 address of the destination peer. This option and `index` provide the same level of output.
- `brief`—(Optional) Display standard information about all existing IKE SAs. (Default)
- `detail`—(Optional) Display detailed information about all existing IKE SAs.
- `family`—(Optional) Display IKE SAs by family. This option is used to filter the output.
  - `inet`—IPv4 address family.
  - `inet6`—IPv6 address family.
- `fpc` *slot-number*—(Optional) Display information about existing IKE SAs in this Flexible PIC Concentrator (FPC) slot. This option is used to filter the output.

  In a chassis cluster, when you execute the CLI command `show security ike security-associations pic` *<slot-number>* `fpc` *<slot-number>* in operational mode, only the primary node information about the existing IPsec SAs in the specified Flexible PIC Concentrator (FPC) slot and PIC slot is displayed.
- `index` *SA-index-number*—(Optional) Display information for a particular SA based on the index number of the SA. For a particular SA, display the list of existing SAs by using the command with no options. This option and *peer-address* provide the same level of output.
- `kmd-instance` —(Optional) Display information about existing IKE SAs in the key management process (in this case, it is KMD) identified by FPC *slot-number* and PIC *slot-number*. This option is used to filter the output.
  - `all`—All KMD instances running on the Services Processing Unit (SPU).
  - *kmd-instance-name*—Name of the KMD instance running on the SPU.

- `pic slot-number` —(Optional) Display information about existing IKE SAs in this PIC slot. This option is used to filter the output.
- `sa-type`—(Optional for ADVPN) Type of SA. `shortcut` is the only option for this release.
- `ha-link-encryption`—(Optional) Display information related to interchassis link tunnel only. See *ipsec (High Availability)* and "show security ike security-associations ha-link-encryption (SRX5400, SRX5600, SRX5800)" on page 1170.
- `srg-id`—(Optional) Display information related to a specific services redundancy group (SRG).

Required Privilege Level

view

Output Fields

Table 69 on page 1155 lists the output fields for the `show security ike security-associations` command. Output fields are listed in the approximate order in which they appear.

Table 69: show security ike security-associations Output Fields

Field Name	Field Description
`IKE Peer or Remote Address`	IP address of the destination peer with which the local peer communicates.
`Index`	Index number of an SA. This number is an internally generated number you can use to display information about a single SA.
`Gateway Name`	Name of the IKE gateway.

Table 69: show security ike security-associations Output Fields *(Continued)*

Field Name	Field Description
Location	• FPC—Flexible PIC Concentrator (FPC) slot number. • PIC—PIC slot number. • KMD-Instance—The name of the KMD instance running on the SPU, identified by FPC *slot-number* and PIC *slot-number*. Currently, 4 KMD instances are running on each SPU, and any particular IKE negotiation is carried out by a single KMD instance.
Role	Part played in the IKE session. The device triggering the IKE negotiation is the initiator, and the device accepting the first IKE exchange packets is the responder.
State	State of the IKE SAs: • DOWN—SA has not been negotiated with the peer. • UP—SA has been negotiated with the peer.
Initiator cookie	Random number, called a cookie, which is sent to the remote node when the IKE negotiation is triggered.
Responder cookie	Random number generated by the remote node and sent back to the initiator as a verification that the packets were received. A cookie is aimed at protecting the computing resources from attack without spending excessive CPU resources to determine the cookie's authenticity.

Table 69: show security ike security-associations Output Fields *(Continued)*

Field Name	Field Description
`Exchange type`	Negotiation method agreed on by the two IPsec endpoints, or peers, used to exchange information between one another. Each exchange type or mode determines the number of messages and the payload types that are contained in each message. The modes are: • `main`—The exchange is done with six messages. This mode encrypts the payload, protecting the identity of the neighbor. • `aggressive`—The exchange is done with three messages. This mode does not encrypt the payload, leaving the identity of the neighbor unprotected. IKEv2 protocol does not use the mode configuration for negotiation. Therefore, the mode displays the version number of the security association.
`Authentication method`	Method used to authenticate the source of IKE messages, which can be either `Pre-shared-keys` or digital certificates, such as `DSA-signatures`, `ECDSA-signatures-256`, `ECDSA-signatures-384`, or `RSA-signatures`.
`Local`	Address of the local peer.
`Remote`	Address of the remote peer.
`Lifetime`	Number of seconds remaining until the IKE SA expires.
`Reauth Lifetime`	When enabled, number of seconds remaining until reauthentication triggers a new IKEv2 SA negotiation.
`IKE Fragmentation`	`Enabled` means that both the IKEv2 initiator and responder support message fragmentation and have negotiated the support during the IKE_SA_INIT message exchange. `Size` shows the maximum size of an IKEv2 message before it is fragmented.

Table 69: show security ike security-associations Output Fields *(Continued)*

Field Name	Field Description

Algorithms	IKE algorithms used to encrypt and secure exchanges between the peers during the IPsec Phase 2 process:
	• Authentication—Type of authentication algorithm used:
	  • sha1—Secure Hash Algorithm 1 authentication.
	  • md5—MD5 authentication.
	• Encryption—Type of encryption algorithm used:
	  • aes-256-cbc—Advanced Encryption Standard (AES) 256-bit encryption.
	  • aes-192-cbc— AES192-bit encryption.
	  • aes-128-cbc—AES 128-bit encryption.
	  • 3des-cbc—3 Data Encryption Standard (DES) encryption.
	  • aes-128-gcm—Advanced Encryption Standard (AES) 256-bit encryption.
	  • des-cbc—DES encryption.
	  Starting in Junos OS Release 19.4R2, when you configure aes-128-gcm or aes-256-gcm as an encryption algorithm at the [edit security ipsec proposalproposal-name] hierarchy level, the authentication algorithm field of the show security ikesecurity-associations detail command displays the same configured encryption algorithm.
	• Pseudo random function—Function that generates highly unpredictable random numbers: hmac-md5 or hmac-sha1.
	• Diffie-Hellman group—Specifies the type of Diffie-Hellman group when performing the new Diffie-Hellman exchange. It can be one of the following:
	  • group1—768-bit Modular Exponential (MODP) algorithm.
	  • group2—1024-bit MODP algorithm.
	  • group14—2048-bit MODP group.

Table 69: show security ike security-associations Output Fields *(Continued)*

Field Name	Field Description
	• group15—3072-bit MODP algorithm. • group16—4096-bit MODP algorithm. • group19—256-bit random Elliptic Curve Groups modulo a prime (ECP group) algorithm. • group20—384-bit random ECP group algorithm. • group21—521-bit random ECP group algorithm. • group24—2048-bit MODP group with 256-bit prime order subgroup.
Traffic statistics	• Input bytes–Number of bytes received. • Output bytes–Number of bytes transmitted. • Input packets–Number of packets received. • Output packets–Number of packets transmitted. • Input fragmented packets—Number of IKEv2 fragmented packets received. • Output fragmented packets—Number of IKEv2 fragmented packets transmitted.
Flags	Notification to the key management process of the status of the IKE negotiation: • caller notification sent—Caller program notified about the completion of the IKE negotiation. • waiting for done—Negotiation is done. The library is waiting for the remote end retransmission timers to expire. • waiting for remove—Negotiation has failed. The library is waiting for the remote end retransmission timers to expire before removing this negotiation. • waiting for policy manager—Negotiation is waiting for a response from the policy manager.

Table 69: show security ike security-associations Output Fields *(Continued)*

Field Name	Field Description
IPSec security associations	• *number* created: The number of SAs created. • *number* deleted: The number of SAs deleted.

Phase 2 negotiations in progress	Number of Phase 2 IKE negotiations in progress and status information:
	• Negotiation type—Type of Phase 2 negotiation. Junos OS currently supports quick mode.
	• Message ID—Unique identifier for a Phase 2 negotiation.
	• Local identity—Identity of the local Phase 2 negotiation. The format is *id-type-name (proto-name:port-number,[0..id-data-len] = iddata-presentation)*.
	• Remote identity—Identity of the remote Phase 2 negotiation. The format is *id-type-name (proto-name:port-number,[0..id-data-len] = iddata-presentation)*.
	• Flags—Notification to the key management process of the status of the IKE negotiation:
	  • caller notification sent—Caller program notified about the completion of the IKE negotiation.
	  • waiting for done—Negotiation is done. The library is waiting for the remote end retransmission timers to expire.
	  • waiting for remove—Negotiation has failed. The library is waiting for the remote end retransmission timers to expire before removing this negotiation.
	  • waiting for policy manager—Negotiation is waiting for a response from the policy manager.

Local gateway interface	Interface name of the local gateway.
Routing instance	Name of the local gateway routing instance.

Table 69: show security ike security-associations Output Fields *(Continued)*

Field Name	Field Description
IPsec Tunnel IDs	Indicates the list of child IPsec tunnel IDs

Sample Output

show security ike security-associations (IPv4)

```
user@host> show security ike security-associations
Index        Remote Address        State        Initiator cookie        Responder cookie    Mode
8        192.168.1.2        UP        3a895f8a9f620198    9040753e66d700bb    Main
Index        Remote Address        State        fInitiator cookie    Responder cookie    Mode
9        192.168.1.3         UP        5ba96hfa9f65067        70890755b65b80b        Main
```

show security ike security-associations (IPv6)

```
user@host> show security ike security-associations
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address
5       UP     e48efd6a444853cf  0d09c59aafb720be  Aggressive     2001:db8::1112
```

show security ike security-associations detail (SRX300, SRX320, SRX340, SRX345, and SRX550HM Devices)

```
user@host> show security ike security-associations detail
IKE peer 192.168.134.245, Index 2577565, Gateway Name: tropic
  Role: Initiator, State: UP
  Initiator cookie: b869b3424513340a, Responder cookie: 4cb3488cb19397c3
  Exchange type: Main, Authentication method: Pre-shared-keys Trusted CA group: xyz_ca_grp
  Local: 192.168.134.241:500, Remote: 192.168.134.245:500
  Local gateway interface: ge-0/0/0
  Routing instance: default
  Lifetime: Expires in 169 seconds
  Peer ike-id: 192.168.134.245
```

```
  AAA assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes-128-gcm
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                 1012
   Output bytes  :                 1196
   Input  packets:                    4
   Output packets:                    5
  Flags: IKE SA is created
  IPSec security associations: 1 created, 0 deleted
  Phase 2 negotiations in progress: 0

    Negotiation type: Quick mode, Role: Initiator, Message ID: 0
    Local: 192.168.134.241:500, Remote: 192.168.134.245:500
    Local identity: 192.168.134.241
    Remote identity: 192.168.134.245
    Flags: IKE SA is created
IPsec SA Rekey CREATE_CHILD_SA exchange stats:
   Initiator stats:                                Responder stats:
    Request Out             : 1                     Request In             : 0
    Response In             : 1                     Response Out           : 0
    No Proposal Chosen In   : 0                     No Proposal Chosen Out : 0
    Invalid KE In           : 0                     Invalid KE Out         : 0
    TS Unacceptable In      : 0                     TS Unacceptable Out    : 0
    Res DH Compute Key Fail : 0                     Res DH Compute Key Fail: 0
    Res Verify SA Fail      : 0
    Res Verify DH Group Fail: 0
    Res Verify TS Fail      : 0
```

show security ike security-associations detail (SRX5400, SRX5600, and SRX5800 Devices)

```
user@host> show security ike security-associations detail
IKE peer 2.0.0.2, Index 2068, Gateway Name: IKE_GW
  Role: Responder, State: DOWN
  Initiator cookie: aa08091f3d4f1fb6, Responder cookie: 08c89a7add5f9332
  Exchange type: IKEv2, Authentication method: Pre-shared-keys
  Local gateway interface: ge-0/0/3
  Routing instance: default
```

```
  Local: 2.0.0.1:500, Remote: 2.0.0.2:500
  Lifetime: Expires in 186 seconds
  Reauth Lifetime: Disabled
  IKE Fragmentation: Enabled, Size: 576
  Remote Access Client Info: Unknown Client
  Peer ike-id: 2.0.0.2
  AAA assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha256-128
   Encryption            : aes128-cbc
   Pseudo random function: hmac-sha256
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                 704
   Output bytes  :                1408
   Input  packets:                   4
   Output packets:                   4
   Input  fragmented packets:       0
   Output fragmented packets:       0
  IPSec security associations: 4 created, 2 deleted
  Phase 2 negotiations in progress: 1
  IPSec Tunnel IDs:  500766, 500767


Negotiation type: Quick mode, Role: Responder, Message ID: 0
Local: 2.0.0.1:500, Remote: 2.0.0.2:500
Local identity: 2.0.0.1
Remote identity: 2.0.0.2
Flags: IKE SA is created

IPsec SA Rekey CREATE_CHILD_SA exchange stats:
   Initiator stats:                                Responder stats:
    Request Out             : 0                     Request In             : 0
    Response In             : 0                     Response Out           : 0
    No Proposal Chosen In   : 0                     No Proposal Chosen Out : 0
    Invalid KE In           : 0                     Invalid KE Out         : 0
    TS Unacceptable In      : 0                     TS Unacceptable Out    : 0
    Res DH Compute Key Fail : 0                     Res DH Compute Key Fail: 0
    Res Verify SA Fail      : 0
    Res Verify DH Group Fail: 0
    Res Verify TS Fail      : 0
```

command-name

The *show security ike stats* topic lists the output fields for the `show security ike security-associations detail` command.

show security ike security-associations family inet6

```
user@host> show security ike security-associations family inet6
  IKE peer 2001:db8:1212::1112, Index 5, Gateway Name: tropic
  Role: Initiator, State: UP
  Initiator cookie: e48efd6a444853cf, Responder cookie: 0d09c59aafb720be
  Exchange type: Aggressive, Authentication method: Pre-shared-keys
  Local: 2001:db8:1212::1111:500, Remote: 2001:db8:1212::1112:500
  Lifetime: Expires in 19518 seconds
  Peer ike-id: not valid
  AAA assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : sha1
   Encryption            : 3des-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                1568
   Output bytes  :                2748
   Input  packets:                   6
   Output packets:                  23
  Flags: Caller notification sent
  IPSec security associations: 5 created, 0 deleted
  Phase 2 negotiations in progress: 1

    Negotiation type: Quick mode, Role: Initiator, Message ID: 2900338624
    Local: 2001:db8:1212::1111:500, Remote: 2001:db8:1212::1112:500
    Local identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
    Remote identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
    Flags: Caller notification sent, Waiting for done
```

show security ike security-associations index 222075191 detail

```
user@host> show security ike security-associations index 222075191 detail
node0:
```

```
-
IKE peer 192.168.1.2, Index 222075191, Gateway Name: ZTH_HUB_GW
  Location: FPC 0, PIC 3, KMD-Instance 2
  Auto Discovery VPN:
   Type: Static, Local Capability: Suggester, Peer Capability: Partner
   Suggester Shortcut Suggestions Statistics:
     Suggestions sent    :    2
     Suggestions accepted:    4
     Suggestions declined:    1
  Role: Responder, State: UP
  Initiator cookie: 7b996b4c310d2424, Responder cookie: 5724c5882a212157
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 192.168.1.1:500, Remote: 192.168.1.2:500
  Lifetime: Expires in 828 seconds
  Peer ike-id: C=US, DC=example, ST=CA, L=Sunnyvale, O=example, OU=engineering, CN=cssvk36-d
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                20474
   Output bytes  :                21091
   Input  packets:                  237
   Output packets:                  237
  IPSec security associations: 2 created, 0 deleted
  Phase 2 negotiations in progress: 1

    Negotiation type: Quick mode, Role: Responder, Message ID: 0
    Local: 192.168.1.1:500, Remote: 192.168.1.2:500
    Local identity: C=US, DC=example, ST=CA, L=Sunnyvale, O=example, OU=engineering, CN=host1
    Remote identity: C=US, DC=example, ST=CA, L=Sunnyvale, O=example, OU=engineering, CN=host2
    Flags: IKE SA is created
```

show security ike security-associations index 788674 detail

```
user@host> show security ike security-associations index 788674 detail
IKE peer 192.168.1.1, Index 788674, Gateway Name: ZTH_SPOKE_GW
  Auto Discovery VPN:
```

```
   Type: Static, Local Capability: Partner, Peer Capability: Suggester
   Partner Shortcut Suggestions Statistics:
     Suggestions received:    2
     Suggestions accepted:    2
     Suggestions declined:    0
  Role: Initiator, State: UP
  Initiator cookie: 7b996b4c310d2424, Responder cookie: 5724c5882a212157
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 192.168.1.2:500, Remote: 192.168.1.1:500
  Lifetime: Expires in 734 seconds
  Peer ike-id: C=US, DC=example, ST=CA, L=Sunnyvale, O=example, OU=engineering, CN=test
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                22535
   Output bytes  :                21918
   Input  packets:                  256
   Output packets:                  256
  IPSec security associations: 2 created, 0 deleted
  Phase 2 negotiations in progress: 1

    Negotiation type: Quick mode, Role: Initiator, Message ID: 0
    Local: 192.168.1.2:500, Remote: 192.168.1.1:500
    Local identity: C=US, DC=example, ST=CA, L=Sunnyvale, O=example, OU=engineering, CN=host1
    Remote identity: C=US, DC=example, ST=CA, L=Sunnyvale, O=example, OU=engineering, CN=host2
    Flags: IKE SA is created
```

show security ike security-associations 192.168.1.2

```
user@host> show security ike security-associations 192.168.1.2
Index     State  Initiator cookie  Responder cookie  Mode Remote Address
   8        UP     3a895f8a9f620198  9040753e66d700bb  Main 192.168.1.2
```

show security ike security-associations fpc 6 pic 1 kmd-instance all (SRX Series Devices)

```
user@host> show security ike security-associations fpc 6 pic 1 kmd-instance all
Index      Remote Address  State  Initiator cookie  Responder cookie  Mode

1728053250 192.168.1.2     UP     fc959afd1070d10b  bdeb7e8c1ea99483  Main
```

show security ike security-associations detail (ADVPN Suggester, Static Tunnel)

```
user@host> show security ike security-associations detail
IKE peer 192.168.0.105, Index 13563297, Gateway Name: zth_hub_gw
  Location: FPC 0, PIC 0, KMD-Instance 1
  Auto Discovery VPN:
 Type: Static, Local Capability: Suggester, Peer Capability: Partner
   Suggester Shortcut Suggestions Statistics:
     Suggestions sent            :  12
     Suggestion response accepted:  12
     Suggestion response declined:   0
  Role: Responder, State: UP
  Initiator cookie: 4d3f4e4b2e75d727, Responder cookie: 81ab914e13cecd21
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 192.168.0.154:500, Remote: 192.168.0.105:500
  Lifetime: Expires in 26429 seconds
  Peer ike-id: DC=example, CN=host02, L=Sunnyvale, ST=CA, C=US
```

show security ike security-associations detail (ADVPN Partner, Static Tunnel)

```
user@host> show security ike security-associations detail
IKE peer 192.168.0.154, Index 4980720, Gateway Name: zth_spoke_gw
  Location: FPC 0, PIC 0, KMD-Instance 1
  Auto Discovery VPN:
 Type: Static, Local Capability: Partner, Peer Capability: Suggester
   Partner Shortcut Suggestions Statistics:
     Suggestions received:  12
     Suggestions accepted:  12
     Suggestions declined:   0
  Role: Initiator, State: UP
  Initiator cookie: 4d3f4e4b2e75d727, Responder cookie: 81ab914e13cecd21
  Exchange type: IKEv2, Authentication method: RSA-signatures
```

```
  Local: 192.168.0.105:500, Remote: 192.168.0.154:500
  Lifetime: Expires in 26252 seconds
  Peer ike-id: DC=example, CN=host01, OU=SBU, O=example, L=Sunnyvale, ST=CA, C=US
```

show security ike security-associations detail (ADVPN Partner, Shortcut)

```
user@host> show security ike security-associations detail
IKE peer 192.168.0.106, Index 4980737, Gateway Name: GW-ADVPN-GT-ADVPN-zth_spoke_vpn-268173323
  Location: FPC 0, PIC 0, KMD-Instance 1
  Auto Discovery VPN:
   Type: Shortcut, Local Capability: Partner, Peer Capability: Partner
  Role: Responder, State: UP
  Initiator cookie: e1ed0c655929debc, Responder cookie: 437de6ed784ba63e
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 192.168.0.105:500, Remote: 192.168.0.106:500
  Lifetime: Expires in 28796 seconds
  Peer ike-id: DC=example, CN=paulyd, L=Sunnyvale, ST=CA, C=US
```

show security ike security-associations sa-type shortcut (ADVPN)

```
user@host> show security ike security-associations sa-type shortcut
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address
4980742 UP     vb56fbe694eaee5b6 064dbccbfa3b2aab  IKEv2          192.168.0.106
```

show security ike security-associations sa-type shortcut detail (ADVPN)

```
user@host> show security ike security-associations sa-type shortcut detail
IKE peer 192.168.0.106, Index 4980742, Gateway Name: GW-ADVPN-GT-ADVPN-zth_spoke_vpn-268173327
  Location: FPC 0, PIC 0, KMD-Instance 1
  Auto Discovery VPN:
   Type: Shortcut, Local Role: Partner, Peer Role: Partner
  Role: Responder, State: UP
```

show security ike security-associations detail (IKEv2 Reauthentication)

```
user@host> show security ike security-associations detail
IKE peer 10.1.2.11, Index 6009224, Gateway Name: GW
  Role: Responder, State: UP
  Initiator cookie: 2c74d14c798a9d70, Responder cookie: 83cbb49bfbcb80cb
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 10.1.1.11:500, Remote: 10.1.2.11:500
  Lifetime: Expires in 173 seconds
  Reauth Lifetime: Expires in 600 seconds
  Peer ike-id: vsrx@example.net
  AAA assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes128-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-2
  Traffic statistics:
   Input  bytes  :                1782
   Output bytes  :                1743
   Input  packets:                   2
```

show security ike security-associations detail (IKEv2 Fragmentation)

```
user@host> show security ike security-associations detail
IKE peer 172.24.23.157, Index 11883008, Gateway Name: routebased_s2s_gw-552_1
  Role: Responder, State: UP
  Initiator cookie: f3255e720f162e3a, Responder cookie: 17555e3ff7451841
  Exchange type: Main, Authentication method: Pre-shared-keys Trusted CA group: xyz_ca_grp
  Local: 192.168.254.1:500, Remote: 172.24.23.157:500
  Lifetime: Expires in 530 seconds
  Reauth Lifetime: Disabled
  IKE Fragmentation: Enabled, Size: 576
  Peer ike-id: 172.24.23.157
  AAA assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : 3des-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
```

```
  Traffic statistics:
   Input  bytes  :                1004
   Output bytes  :                 756
   Input  packets:                   6
   Output packets:                   4
   Input  fragmented packets:  3
   Output fragmented packets: 3
  IPSec security associations: 1 created, 1 deleted
  Phase 2 negotiations in progress: 1

    Negotiation type: Quick mode, Role: Responder, Message ID: 0
    Local: 192.168.254.1:500, Remote: 172.24.23.157:500
    Local identity: 192.168.254.1
    Remote identity: 172.24.23.157
    Flags: IKE SA is created
```

show security ike security-associations ha-link-encryption (SRX5400, SRX5600, SRX5800)

Starting in Junos OS Release 20.4R1, when you configure the high availability (HA) feature, you can use this show command to view only interchassis link tunnel details. The following command displays only the link encryption SAs on both nodes.

```
user@host> show security ike security-associations ha-link-encryption

Index      State  Initiator cookie    Responder cookie  Mode   Remote Address
4294966287  UP    7b77b4e2fd5a87e5    ab4a398e6a28687a  IKEv2   23.0.0.2
```

show security ike security-associations srg-id

```
user@host> show security ike security-associations srg-id 1
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address
16778113 UP    16d1f4efae91608c  53f234767bdd0b9b  IKEv2          10.112.0.1
```

Release Information

Command introduced in Junos OS Release 8.5. Support for the `fpc`, `pic`, and `kmd-instance` options added in Junos OS Release 9.3. Support for the `family` option added in Junos OS Release 11.1. Support for Auto

Discovery VPN added in Junos OS Release 12.3X48-D10. Support for IKEv2 reauthentication added in Junos OS Release 15.1X49-D60. Support for IKEv2 fragmentation added in Junos OS Release 15.1X49-D80.

Support for the `ha-link-encryption` option added in Junos OS Release 20.4R1.

Support for the `srg-id` option added in Junos OS Release 22.4R1.

RELATED DOCUMENTATION

Example: Configuring a Route-Based VPN Tunnel in a User Logical Systems

show security ipsec security-associations

IN THIS SECTION

- Syntax | **1171**
- Description | **1172**
- Options | **1172**
- Required Privilege Level | **1173**
- Output Fields | **1173**
- Sample Output | **1184**
- show security ipsec security-associations detail (SRX Series devices and MX Series Routers) | **1205**
- Release Information | **1208**

Syntax

```
show security ipsec security-associations
<brief | detail>
<family (inet  | inet6)>
<fpc slot-number pic slot-number>
<index SA-index-number>
<kmd-instance (all | kmd-instance-name)>
```

```
<pic slot-number fpc slot-number>
<sa-type shortcut>
<traffic-selector traffic-selector-name>
<srg-id id-number>
<vpn-name vpn-name>
<ha-link-encryption>
```

Description

Display information about the IPsec security associations (SAs).

In Junos OS Releases 20.1R2, 20.2R2, 20.3R2, 20.3R1, and later, when you execute the `show security ipsec security-associations detail` command, a new output field `IKE SA Index` corresponding to every IPsec SA within a tunnel is displayed under each IPsec SA information. See "show security ipsec security-associations detail (SRX5400, SRX5600, SRX5800)" on page 1199.

Options

none — Display information about all SAs.

brief | detail — (Optional) Display the specified level of output. The default is `brief`.

family — (Optional) Display SAs by family. This option is used to filter the output.

- `inet`—IPv4 address family.
- `inet6`—IPv6 address family.

fpc *slot-number* pic *slot-number* — (Optional) Display information about existing IPsec SAs in the specified Flexible PIC Concentrator (FPC) slot and PIC slot.

In a chassis cluster, when you execute the CLI command `show security ipsec security-associations pic <slot-number> fpc <slot-number>` in operational mode, only the primary node information about the existing IPsec SAs in the specified Flexible PIC Concentrator (FPC) slot and PIC slot is displayed.

index *SA-index-number* — (Optional) Display detailed information about the specified SA identified by this index number. To obtain a list of all SAs that includes their index numbers, use the command with no options.

kmd-instance	(Optional) Display information about existing IPsec SAs in the key management process (in this case, it is KMD) identified by the FPC *slot-number* and PIC *slot-number*. • `all`—All KMD instances running on the Services Processing Unit (SPU). • *kmd-instance-name*—Name of the KMD instance running on the SPU.
pic *slot-number* **fpc** *slot-number*	(Optional) Display information about existing IPsec SAs in the specified PIC slot and FPC slot.
sa-type	(Optional for ADVPN) Display information for the specified type of SA. `shortcut` is the only option for this release.
traffic-selector *traffic-selector-name*	(Optional) Display information about the specified traffic selector.
vpn-name *vpn-name*	(Optional) Display information about the specified VPN.
ha-link-encryption	(Optional) Display information related to interchassis link tunnel only. See *ipsec (High Availability)*, "show security ipsec security-associations ha-link-encryption (SRX5400, SRX5600, SRX5800)" on page 1200, and "show security ipsec sa detail ha-link-encryption (SRX5400, SRX5600, SRX5800)" on page 1201.
srg-id	(Optional) Display information related to a specific services redundancy group (SRG) in a Multinode High Availability setup.

Required Privilege Level

view

Output Fields

Table 70 on page 1174 lists the output fields for the `show security ipsec security-associations` command, Table 71 on page 1179 lists the output fields for the `show security ipsec sa` command and Table 72 on page 1181. lists the output fields for the `show security ipsec sa detail`. Output fields are listed in the approximate order in which they appear.

Table 70: show security ipsec security-associations

Field Name	Field Description	Level of Output
`Total active tunnels`	Total number of active IPsec tunnels.	`brief`
`ID`	Index number of the SA. You can use this number to get additional information about the SA.	All levels
`Algorithm`	Cryptography used to secure exchanges between peers during the IKE negotiations includes: • An authentication algorithm used to authenticate exchanges between the peers. • An encryption algorithm used to encrypt data traffic.	`brief`
`SPI`	Security parameter index (SPI) identifier. An SA is uniquely identified by an SPI. Each entry includes the name of the VPN, the remote gateway address, the SPIs for each direction, the encryption and authentication algorithms, and keys. The peer gateways each have two SAs, one resulting from each of the two phases of negotiation: IKE and IPsec.	`brief`
`Life: sec/kb`	The lifetime of the SA, after which it expires, expressed either in seconds or kilobytes.	`brief`
`Mon`	The Mon field refers to VPN monitoring status. If VPN monitoring is enabled, then this field displays `U` (up) or `D` (down). A hyphen (-) means VPN monitoring is not enabled for this SA. A `V` means that IPsec datapath verification is in progress.	`brief`

Table 70: show security ipsec security-associations *(Continued)*

Field Name	Field Description	Level of Output
`lsys`	The root system.	`brief`
`Port`	If Network Address Translation (NAT) is used, this value is 4500. Otherwise, it is the standard IKE port, 500.	All levels
`Gateway`	IP address of the remote gateway.	`brief`
`Virtual-system`	Name of the logical system.	`detail`
`VPN name`	IPsec name for VPN.	`detail`
`State`	State has two options, `Installed` and `Not Installed`. • `Installed`—The SA is installed in the SA database. • `Not Installed`—The SA is not installed in the SA database. For transport mode, the value of State is always `Installed`.	`detail`
`Local gateway`	Gateway address of the local system.	`detail`
`Remote gateway`	Gateway address of the remote system.	`detail`
`Traffic selector`	Name of the traffic selector.	`detail`
`Local identity`	Identity of the local peer so that its partner destination gateway can communicate with it. The value is specified as an IP address, fully qualified domain name, e-mail address, or distinguished name (DN).	`detail`

Table 70: show security ipsec security-associations *(Continued)*

Field Name	Field Description	Level of Output
`Remote identity`	IP address of the destination peer gateway.	`detail`
`Term`	Defines local IP range, remote IP range, source port range, destination port range, and protocol.	`detail`
`Source-port`	Source port range configured for a term.	`detail`
`Destination-Port`	Destination port range configured for a term.	`detail`
`Version`	IKE version, either `IKEv1` or `IKEv2`.	`detail`
`DF-bit`	State of the don't fragment bit: `set` or `cleared`.	`detail`
`Location`	`FPC`—Flexible PIC Concentrator (FPC) slot number. `PIC`—PIC slot number. `KMD-Instance`—The name of the KMD instance running on the SPU, identified by FPC *slot-number* and PIC *slot-number*. Currently, 4 KMD instances running on each SPU, and any particular IPsec negotiation is carried out by a single KMD instance.	`detail`
`Tunnel events`	Tunnel event and the number of times the event has occurred. See Tunnel Events for descriptions of tunnel events and the action you can take.	`detail`
`Anchorship`	Anchor thread ID for the SA (for SRX4600 Series devices with the `detail` option).	

Table 70: show security ipsec security-associations *(Continued)*

Field Name	Field Description	Level of Output
Direction	Direction of the SA; it can be inbound or outbound.	detail
AUX-SPI	Value of the auxiliary security parameter index(SPI). • When the value is AH or ESP, AUX-SPI is always 0. • When the value is AH+ESP, AUX-SPI is always a positive integer.	detail
Mode	Mode of the SA: • transport—Protects host-to-host connections. • tunnel–Protects connections between security gateways.	detail
Type	Type of the SA: • manual—Security parameters require no negotiation. They are static and are configured by the user. • dynamic—Security parameters are negotiated by the IKE protocol. Dynamic SAs are not supported in transport mode.	detail
State	State of the SA: • Installed—The SA is installed in the SA database. • Not Installed—The SA is not installed in the SA database. For transport mode, the value of State is always Installed.	detail

Table 70: show security ipsec security-associations *(Continued)*

Field Name	Field Description	Level of Output
`Protocol`	Protocol supported. • Transport mode supports Encapsulation Security Protocol (ESP) and Authentication Header (AH). • Tunnel mode supports ESP and AH.	`detail`
`Authentication`	Type of authentication used.	`detail`
`Encryption`	Type of encryption used. Starting in Junos OS Release 19.4R2, when you configure `aes-128-gcm` or `aes-256-gcm` as an encryption algorithm at the `[edit security ipsec proposal proposal-name]` hierarchy level, the authentication algorithm field of the `show security ipsec security-associations detail` command displays the same configured encryption algorithm.	`detail`
`Soft lifetime`	The soft lifetime informs the IPsec key management system that the SA is about to expire. Each lifetime of an SA has two display options, hard and soft, one of which must be present for a dynamic SA. This allows the key management system to negotiate a new SA before the hard lifetime expires. • `Expires in seconds`—Number of seconds left until the SA expires.	`detail`
`Hard lifetime`	The hard lifetime specifies the lifetime of the SA. • `Expires in seconds`—Number of seconds left until the SA expires.	`detail`

Table 70: show security ipsec security-associations *(Continued)*

Field Name	Field Description	Level of Output
Lifesize Remaining	The lifesize remaining specifies the usage limits in kilobytes. If there is no lifesize specified, it shows unlimited. • Expires in kilobytes—Number of kilobytes left until the SA expires.	detail
Anti-replay service	State of the service that prevents packets from being replayed. It can be Enabled or Disabled.	detail
Replay window size	Size of the antireplay service window, which is 64 bits.	detail
Bind-interface	The tunnel interface to which the route-based VPN is bound.	detail
Copy-Outer-DSCP	Indicates if the system copies the outer DSCP value from the IP header to the inner IP header.	detail
tunnel-establishment	Indicates how the IKE is activated.	detail
IKE SA index	Indicates the list of parent IKE security associations.	detail

Table 71: show security ipsec sa Output Fields

Field Name	Field Description
Total active tunnels	Total number of active IPsec tunnels.

Table 71: show security ipsec sa Output Fields *(Continued)*

Field Name	Field Description
ID	Index number of the SA. You can use this number to get additional information about the SA.
Algorithm	Cryptography used to secure exchanges between peers during the IKE Phase 2 negotiations includes: • An authentication algorithm used to authenticate exchanges between the peers. Options are hmac-md5-96, hmac-sha-256-128, or hmac-sha1-96. • An encryption algorithm used to encrypt data traffic. Options are 3des-cbc, aes-128-cbc, aes-192-cbc, aes-256-cbc, or des-cbc.
SPI	Security parameter index (SPI) identifier. An SA is uniquely identified by an SPI. Each entry includes the name of the VPN, the remote gateway address, the SPIs for each direction, the encryption and authentication algorithms, and keys. The peer gateways each have two SAs, one resulting from each of the two phases of negotiation: Phase 1 and Phase 2.
Life:sec/kb	The lifetime of the SA, after which it expires, expressed either in seconds or kilobytes.
Mon	The Mon field refers to VPN monitoring status. If VPN monitoring is enabled, then this field displays U (up) or D (down). A hyphen (-) means VPN monitoring is not enabled for this SA. A V means that IPSec datapath verification is in progress.
lsys	The root system.
Port	If Network Address Translation (NAT) is used, this value is 4500. Otherwise, it is the standard IKE port, 500.
Gateway	Gateway address of the system.

Table 72: show security ipsec sa detail Output Fields

Field Name	Field Description
ID	Index number of the SA. You can use this number to get additional information about the SA.
Virtual-system	The virtual system name.
VPN Name	IPSec name for VPN.
Local Gateway	Gateway address of the local system.
Remote Gateway	Gateway address of the remote system.
Local Identity	Identity of the local peer so that its partner destination gateway can communicate with it. The value is specified as an IP address, fully qualified domain name, e-mail address, or distinguished name (DN).
Remote Identity	IP address of the destination peer gateway.
Version	IKE version. For example, IKEv1, IKEv2.
Passive Mode Tunneling	IPsec tunneling of malformed packets; enabled if set or disabled if not set.
DF-bit	State of the don't fragment bit: set or cleared.
Bind-interface	The tunnel interface to which the route-based VPN is bound.
Tunnel Events	
Direction	Direction of the SA; it can be inbound or outbound.

Table 72: show security ipsec sa detail Output Fields *(Continued)*

Field Name	Field Description
AUX-SPI	Value of the auxiliary security parameter index(SPI). • When the value is AH or ESP, AUX-SPI is always 0. • When the value is AH+ESP, AUX-SPI is always a positive integer.
VPN Monitoring	If VPN monitoring is enabled, then the Mon field displays U (up) or D (down). A hyphen (-) means VPN monitoring is not enabled for this SA. A V means that IPsec datapath verification is in progress.
Hard lifetime	The hard lifetime specifies the lifetime of the SA. • Expires in seconds - Number of seconds left until the SA expires.
Lifesize Remaining	The lifesize remaining specifies the usage limits in kilobytes. If there is no lifesize specified, it shows unlimited.
Soft lifetime	The soft lifetime informs the IPsec key management system that the SA is about to expire. Each lifetime of an SA has two display options, hard and soft, one of which must be present for a dynamic SA. This allows the key management system to negotiate a new SA before the hard lifetime expires. • Expires in seconds - Number of seconds left until the SA expires.
Mode	Mode of the SA: • transport - Protects host-to-host connections. • tunnel - Protects connections between security gateways.
Type	Type of the SA: • manual - Security parameters require no negotiation. They are static and are configured by the user. • dynamic - Security parameters are negotiated by the IKE protocol. Dynamic SAs are not supported in transport mode.

Table 72: show security ipsec sa detail Output Fields *(Continued)*

Field Name	Field Description
`State`	State of the SA: • `Installed` - The SA is installed in the SA database. • `Not Installed` - The SA is not installed in the SA database. For transport mode, the value of State is always Installed.

`Protocol`	Protocol supported.
	• Transport mode supports Encapsulation Security Protocol (ESP) and Authentication Header (AH).
	• Tunnel mode supports ESP and AH.
	  • `Authentication` - Type of authentication used.
	  • `Encryption` - Type of encryption used.

`Anti-replay service`	State of the service that prevents packets from being replayed. It can be `Enabled` or `Disabled`.
`Replay window size`	Configured size of the antireplay service window. It can be 32 or 64 packets. If the replay window size is 0, the antireplay service is disabled. The antireplay window size protects the receiver against replay attacks by rejecting old or duplicate packets.
Interchassis Link Tunnel	
HA Link Encryption Mode	High availability mode supported. Displays `Multi-Node` when multi-node high availability feature is enabled.

Sample Output

For brevity, the show command outputs does not display all the values of the configuration. Only a subset of the configuration is displayed. Rest of the configuration on the system has been replaced with ellipses (...).

show security ipsec security-associations (IPv4)

```
user@host> show security ipsec security-associations
  Total active tunnels: 14743 Total Ipsec sas: 14743
  ID    Algorithm       SPI      Life:sec/kb  Mon lsys Port  Gateway
  <511672 ESP:aes-cbc-128/sha1 0x071b8cd2      -   root 500   10.21.45.152
  >503327 ESP:aes-cbc-128/sha1 0x69d364dd 1584/ unlim - root 500 10.21.12.255
  <503327 ESP:aes-cbc-128/sha1 0x0a577f2d 1584/ unlim - root 500 10.21.12.255
  >512896 ESP:aes-cbc-128/sha1 0xd2f51c81 1669/ unlim - root 500 10.21.50.96
  <512896 ESP:aes-cbc-128/sha1 0x071b8d9e 1669/ unlim - root 500 10.21.50.96
  >513881 ESP:aes-cbc-128/sha1 0x95955834 1696/ unlim - root 500 10.21.54.57
  <513881 ESP:aes-cbc-128/sha1 0x0a57860c 1696/ unlim - root 500 10.21.54.57
  >505835 ESP:aes-cbc-128/sha1 0xf827b5c6 1598/ unlim - root 500 10.21.22.204
  <505835 ESP:aes-cbc-128/sha1 0x0f43bf3f 1598/ unlim - root 500 10.21.22.204
  >506531 ESP:aes-cbc-128/sha1 0x01694572 1602/ unlim - root 500 10.21.25.131
  <506531 ESP:aes-cbc-128/sha1 0x0a578143 1602/ unlim - root 500 10.21.25.131
  >512802 ESP:aes-cbc-128/sha1 0xdc292de4 1668/ unlim - root 500 10.21.50.1
  <512802 ESP:aes-cbc-128/sha1 0x0a578558 1668/ unlim - root 500 10.21.50.1
  >512413 ESP:aes-cbc-128/sha1 0xbe2c52d5 1660/ unlim - root 500 10.21.48.125
  <512413 ESP:aes-cbc-128/sha1 0x1129580c 1660/ unlim - root 500 10.21.48.125
  >505075 ESP:aes-cbc-128/sha1 0x2aae6647 1593/ unlim - root 500 10.21.19.213
  <505075 ESP:aes-cbc-128/sha1 0x02dc5c50 1593/ unlim - root 500 10.21.19.213
  >514055 ESP:aes-cbc-128/sha1 0x2b8adfcb 1704/ unlim - root 500 10.21.54.238
  <514055 ESP:aes-cbc-128/sha1 0x0f43c49a 1704/ unlim - root 500 10.21.54.238
  >508898 ESP:aes-cbc-128/sha1 0xbcced4d6 1619/ unlim - root 500 10.21.34.194
  <508898 ESP:aes-cbc-128/sha1 0x1492035a 1619/ unlim - root 500 10.21.34.194
  >505328 ESP:aes-cbc-128/sha1 0x2a8d2b36 1594/ unlim - root 500 10.21.20.208
  <505328 ESP:aes-cbc-128/sha1 0x14920107 1594/ unlim - root 500 10.21.20.208
  >500815 ESP:aes-cbc-128/sha1 0xdd86c89a 1573/ unlim - root 500 10.21.3.47
  <500815 ESP:aes-cbc-128/sha1 0x1129507f 1573/ unlim - root 500 10.21.3.47
  >503758 ESP:aes-cbc-128/sha1 0x64cc490e 1586/ unlim - root 500 10.21.14.172
  <503758 ESP:aes-cbc-128/sha1 0x14920001 1586/ unlim - root 500 10.21.14.172
  >504004 ESP:aes-cbc-128/sha1 0xde0b63ee 1587/ unlim - root 500 10.21.15.164
  <504004 ESP:aes-cbc-128/sha1 0x071b87d4 1587/ unlim - root 500 10.21.15.164
  >508816 ESP:aes-cbc-128/sha1 0x2703b7a5 1618/ unlim - root 500 10.21.34.112
```

```
  <508816 ESP:aes-cbc-128/sha1 0x071b8af6 1618/ unlim - root 500 10.21.34.112
  >511341 ESP:aes-cbc-128/sha1 0x828f3330 1644/ unlim - root 500 10.21.44.77
  <511341 ESP:aes-cbc-128/sha1 0x02dc6064 1644/ unlim - root 500 10.21.44.77
  >500456 ESP:aes-cbc-128/sha1 0xa6f1515d 1572/ unlim - root 500 10.21.1.200
  <500456 ESP:aes-cbc-128/sha1 0x1491fddb 1572/ unlim - root 500 10.21.1.200
  >512506 ESP:aes-cbc-128/sha1 0x4108f3a3 1662/ unlim - root 500 10.21.48.218
  <512506 ESP:aes-cbc-128/sha1 0x071b8d5d 1662/ unlim - root 500 10.21.48.218
  >504657 ESP:aes-cbc-128/sha1 0x27a6b8b3 1591/ unlim - root 500 10.21.18.41
  <504657 ESP:aes-cbc-128/sha1 0x112952fe 1591/ unlim - root 500 10.21.18.41
  >506755 ESP:aes-cbc-128/sha1 0xc0afcff0 1604/ unlim - root 500 10.21.26.100
  <506755 ESP:aes-cbc-128/sha1 0x149201f5 1604/ unlim - root 500 10.21.26.100
  >508023 ESP:aes-cbc-128/sha1 0xa1a90af8 1612/ unlim - root 500 10.21.31.87
  <508023 ESP:aes-cbc-128/sha1 0x02dc5e3b 1612/ unlim - root 500 10.21.31.87
  >509190 ESP:aes-cbc-128/sha1 0xee52074d 1621/ unlim - root 500 10.21.35.230
  <509190 ESP:aes-cbc-128/sha1 0x0f43c16e 1621/ unlim - root 500 10.21.35.230
  >505051 ESP:aes-cbc-128/sha1 0x24130b1c 1593/ unlim - root 500 10.21.19.188
  <505051 ESP:aes-cbc-128/sha1 0x149200d9 1593/ unlim - root 500 10.21.19.188
  >513214 ESP:aes-cbc-128/sha1 0x2c4752d1 1676/ unlim - root 500 10.21.51.158
  <513214 ESP:aes-cbc-128/sha1 0x071b8dd3 1676/ unlim - root 500 10.21.0.51.158
  >510808 ESP:aes-cbc-128/sha1 0x4acd94d3 1637/ unlim - root 500 10.21.42.56
  <510808 ESP:aes-cbc-128/sha1 0x071b8c42 1637/ unlim - root 500 10.21.42.56
```

show security ipsec security-associations (IPv6)

```
user@host> show security ipsec security-associations
Total active tunnels: 1
ID     Algorithm         SPI      Life:sec/kb  Mon  vsys Port  Gateway
131074 ESP:aes256/sha256 14caf1d9 3597/ unlim   -   root 500   2001:db8::1112
131074 ESP:aes256/sha256 9a4db486 3597/ unlim   -   root 500   2001:db8::1112
```

show security ipsec security-associations index 511672

```
user@host> show security ipsec security-associations index 511672
ID: 511672 Virtual-system: root, VPN Name: ipsec_vpn
  Local Gateway: 10.20.0.1, Remote Gateway: 10.21.45.152
  Traffic Selector Name: ts
  Local Identity: ipv4(10.191.151.0-10.191.151.255)
  Remote Identity: ipv4(10.40.151.0-10.40.151.255)
  Version: IKEv2
  DF-bit: clear, Copy-Outer-DSCP Disabled, Bind-interface: st0.0, Policy-name: IPSEC_POL
```

```
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0
  Multi-sa, Configured SAs# 0, Negotiated SAs#: 0
  Location: FPC 0, PIC 1, KMD-Instance 0
  Anchorship: Thread 10
  Direction: inbound, SPI: 0x835b8b42, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 1639 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1257 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
  Direction: outbound, SPI: 0x071b8cd2, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 1639 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1257 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
```

show security ipsec security-associations index 131073 detail

```
user@host> show security ipsec security-associations index 131073 detail
ID: 131073 Virtual-system: root, VPN Name: IPSEC_VPN1
  Local Gateway: 10.4.0.1, Remote Gateway: 10.5.0.1
  Local Identity: ipv4_subnet(any:0,[0..7]=10.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=10.0.0.0/0)
  Version: IKEv2
  DF-bit: clear, Copy-Outer-DSCP Disabled, Bind-interface: st0.1
  Port: 500, Nego#: 18, Fail#: 0, Def-Del#: 0 Flag: 0x600a39
  Multi-sa, Configured SAs# 9, Negotiated SAs#: 9
  Tunnel events:
    Mon Apr 23 2018 22:20:54 -0700: IPSec SA negotiation successfully completed (1 times)
    Mon Apr 23 2018 22:20:54 -0700: IKE SA negotiation successfully completed (2 times)
    Mon Apr 23 2018 22:20:18 -0700: User cleared IKE SA from CLI, corresponding IPSec SAs
cleared (1 times)
    Mon Apr 23 2018 22:19:55 -0700: IPSec SA negotiation successfully completed (2 times)
    Mon Apr 23 2018 22:19:23 -0700: Tunnel is ready. Waiting for trigger event or peer to
trigger negotiation (1 times)
    Mon Apr 23 2018 22:19:23 -0700: Bind-interface's zone received. Information updated (1 times)
```

```
    Mon Apr 23 2018 22:19:23 -0700: External interface's zone received. Information updated (1
times)
  Direction: inbound, SPI: 2d8e710b, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 1930 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1563 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes256-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64
    Multi-sa FC Name: default
  Direction: outbound, SPI: 5f3a3239, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 1930 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1563 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes-256-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64
    Multi-sa FC Name: default
  Direction: inbound, SPI: 5d227e19, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 1930 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1551 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes-256-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64
    Multi-sa FC Name: best-effort
  Direction: outbound, SPI: 5490da, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 1930 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1551 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes-256-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64
...
```

Starting with Junos OS Release 18.2R1, the CLI `show security ipsec security-associations index` *`index-number`* `detail` output displays all the child SA details including forwarding class name.

show security ipsec sa

```
user@host> show security ipsec sa
Total active tunnels: 2
ID Algorithm SPI Life:sec/kb Mon lsys Port Gateway
>67108885 ESP:aes-gcm-256/None fdef4dab 2918/ unlim - root 500 2001:db8:3000::2
>67108885 ESP:aes-gcm-256/None e785dadc 2918/ unlim - root 500 2001:db8:3000::2
>67108887 ESP:aes-gcm-256/None 34a787af 2971/ unlim - root 500 2001:db8:5000::2
>67108887 ESP:aes-gcm-256/None cf57007f 2971/ unlim - root 500 2001:db8:5000::2
```

show security ipsec sa detail

```
user@host> show security ipsec sa detail
ID: 500201 Virtual-system: root, VPN Name: IPSEC_VPN
  Local Gateway: 10.2.0.1, Remote Gateway: 10.2.0.2
  Local Identity: ipv4(10.0.0.0-255.255.255.255)
  Remote Identity: ipv4(10.0.0.0-255.255.255.255)
  Version: IKEv1
  DF-bit: clear, Copy-Outer-DSCP Disabled, Bind-interface: st0.1, Policy-name: IPSEC_POL
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0
  Multi-sa, Configured SAs# 0, Negotiated SAs#: 0
  Location: FPC 0, PIC 1, KMD-Instance 0
  Anchorship: Thread 1
  Distribution-Profile: default-profile
  Direction: inbound, SPI: 0x0a25c960, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 91 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 44 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64
    tunnel-establishment: establish-tunnels-responder-only-no-rekey
  Direction: outbound, SPI: 0x43e34ad3, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 91 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 44 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
```

```
    Anti-replay service: counter-based enabled, Replay window size: 64
    tunnel-establishment: establish-tunnels-responder-only-no-rekey
...
```

Starting with Junos OS Release 19.1R1, a new field **tunnel-establishment** in the output of the CLI `show security ipsec sa detail` displays the option configured under `ipsec vpn establish-tunnels` hierarchy.

Starting with Junos OS Release 21.3R1, a new field **Tunnel MTU** in the output of the CLI `show security ipsec sa detail` displays the option configured under `ipsec vpn hub-to-spoke-vpn tunnel-mtu` hierarchy.

Starting in Junos OS Release 22.1R3, on SRX5000 line of devices, the Tunnel MTU is not displayed in the CLI output if the tunnel MTU is not configured.

show security ipsec sa details (MX-SPC3)

```
user@host>show security ipsec sa detailID: 500055 Virtual-system: root, VPN Name: IPSEC_VPN
  Local Gateway: 10.2.0.1, Remote Gateway: 10.2.0.2
  Local Identity: ipv4(10.0.0.0-255.255.255.255)
  Remote Identity: ipv4(10.0.0.0-255.255.255.255)
  Version: IKEv2
  DF-bit: clear, Copy-Outer-DSCP Disabled, Bind-interface: st0.1, Tunnel MTU: 1420  Policy-name:
IPSEC_POL
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0
  Multi-sa, Configured SAs# 0, Negotiated SAs#: 0
  Location: FPC 0, PIC 0, KMD-Instance 0
  Anchorship: Thread 15
  Distribution-Profile: default-profile
  Direction: inbound, SPI: 0x229b998e, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 23904 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 23288 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-md5-96, Encryption: aes-cbc (128 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
    Extended-Sequence-Number: Enabled
    tunnel-establishment: establish-tunnels-immediately
  Direction: outbound, SPI: 0xb2e843a3, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 23904 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 23288 seconds
```

```
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-md5-96, Encryption: aes-cbc (128 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
    Extended-Sequence-Number: Enabled
    tunnel-establishment: establish-tunnels-immediately
```

show security ipsec sa details (MX-SPC3) with passive mode tunneling

```
user@host>show security ipsec sa detail
ID: 500054 Virtual-system: root, VPN Name: TUN_3
  Local Gateway: 100.0.0.3, Remote Gateway: 200.0.0.3
  Traffic Selector Name: ts1
  Local Identity: ipv4(11.0.0.3-11.0.0.3)
  Remote Identity: ipv4(75.0.0.3-75.0.0.3)
  TS Type: traffic-selector
  Version: IKEv2
  Quantum Secured: No
  PFS group: N/A
  SRG ID: 0
  Passive mode tunneling: Enabled
  DF-bit: clear, Copy-Outer-DSCP Disabled, Bind-interface: st0.3, Policy-name: IPSEC_POLICY
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0
  Multi-sa, Configured SAs# 0, Negotiated SAs#: 0
  Tunnel events:
    Mon Sep 19 2022 19:27:44: IPsec SA negotiation succeeds (1 times)
  Location: FPC 3, PIC 1, KMD-Instance 0
  Anchorship: Thread 15
  Distribution-Profile: vms-3/1/0
  Direction: inbound, SPI: 0x25c03740, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expired
    Lifesize Remaining: Expired
    Soft lifetime: Expires in 2920 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: aes256-gcm, Encryption: aes-gcm (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 512
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-immediately
    IKE SA Index: 122
  Direction: outbound, SPI: 0x8e8f2009, AUX-SPI: 0
                              , VPN Monitoring: -
```

```
    Hard lifetime: Expired
    Lifesize Remaining: Expired
    Soft lifetime: Expires in 2920 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: aes256-gcm, Encryption: aes-gcm (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 512
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-immediately
    IKE SA Index: 122
```

show security ipsec security-association

```
user@host>show security ipsec security-association
Total active tunnels: 1     Total IPsec sas: 1
  ID      Algorithm       SPI      Life:sec/kb  Mon lsys Port  Gateway
  <500006 ESP:aes-gcm-128/aes128-gcm 0x782b233c 1432/ unlim - root 500 10.2.0.2
```

show security ipsec security-associations brief

```
user@host> show security ipsec security-associations brief
Total active tunnels: 2     Total Ipsec sas: 18
  ID    Algorithm       SPI      Life:sec/kb  Mon lsys Port  Gateway
  <131073 ESP:aes256/sha256 89e5098  1569/ unlim   -   root 500   10.5.0.1
  >131073 ESP:aes256/sha256 fcee9d54 1569/ unlim   -   root 500   10.5.0.1
  <131073 ESP:aes256/sha256 f3117676 1609/ unlim   -   root 500   10.5.0.1
  >131073 ESP:aes256/sha256 6050109f 1609/ unlim   -   root 500   10.5.0.1
  <131073 ESP:aes256/sha256 e01f54b1 1613/ unlim   -   root 500   10.5.0.1
  >131073 ESP:aes256/sha256 29a05dd6 1613/ unlim   -   root 500   10.5.0.1
  <131073 ESP:aes256/sha256 606c90f6 1616/ unlim   -   root 500   10.5.0.1
  >131073 ESP:aes256/sha256 9b5b059d 1616/ unlim   -   root 500   10.5.0.1
  <131073 ESP:aes256/sha256 b8116d6d 1619/ unlim   -   root 500   10.5.0.1
  >131073 ESP:aes256/sha256 b7ed6bfd 1619/ unlim   -   root 500   10.5.0.1
  <131073 ESP:aes256/sha256 4f5ce754 1619/ unlim   -   root 500   10.5.0.1
  >131073 ESP:aes256/sha256 af8984b6 1619/ unlim   -   root 500   10.5.0.1
...
```

show security ipsec security-associations detail

```
user@host> show security ipsec security-associations detail

ID: 500009 Virtual-system: root, VPN Name: IPSEC_VPN
  Local Gateway: 10.2.0.2, Remote Gateway: 10.2.0.1
  Local Identity: ipv4(10.0.0.0-255.255.255.255)
  Remote Identity: ipv4(10.0.0.0-255.255.255.255)
  Version: IKEv1
  PFS group: DH-group-14
  DF-bit: clear, Copy-Outer-DSCP Disabled, Bind-interface: st0.1, Policy-name: IPSEC_POL
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0
  Multi-sa, Configured SAs# 0, Negotiated SAs#: 0
  Location: FPC 0, PIC 0, KMD-Instance 0
  Anchorship: Thread 0
  Distribution-Profile: default-profile
  IKE SA Index: 2068
  Direction: inbound, SPI: 0xba7bb1f2, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 146 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 101 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: des-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-on-traffic
  Direction: outbound, SPI: 0x41650a1b, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 146 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 101 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: des-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-on-traffic
```

show security ipsec security-associations family inet6

```
user@host> show security ipsec security-associations family inet6
  Virtual-system: root
  Local Gateway: 2001:db8:1212::1111, Remote Gateway: 2001:db8:1212::1112
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
    DF-bit: clear
    Direction: inbound, SPI: 14caf1d9, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 3440 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2813 seconds
    Mode: tunnel, Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes256-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64

    Direction: outbound, SPI: 9a4db486, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 3440 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2813 seconds
    Mode: tunnel, Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes256-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64
```

show security ipsec security-associations fpc 6 pic 1 kmd-instance all (SRX Series Devices)

```
user@host> show security ipsec security-associations fpc 6 pic 1 kmd-instance all
  Total active tunnels: 1

ID    Gateway          Port  Algorithm           SPI      Life:sec/kb  Mon vsys

<2    192.168.1.2      500   ESP:aes256/sha256   67a7d25d 28280/unlim   -   0

>2    192.168.1.2      500   ESP:aes256/sha256   a23cbcdc 28280/unlim   -   0
```

show security ipsec security-associations detail (ADVPN Suggester, Static Tunnel)

```
user@host> show security ipsec security-associations detail
ID: 70516737 Virtual-system: root, VPN Name: ZTH_HUB_VPN
  Local Gateway: 192.168.1.1, Remote Gateway: 192.168.1.2
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv2
  DF-bit: clear
  Bind-interface: st0.1

  Port: 500, Nego#: 5, Fail#: 0, Def-Del#: 0 Flag: 0x608a29
  Tunnel events:
  Tue Nov 03 2015 01:24:27 -0800: IPSec SA negotiation successfully completed (1 times)
  Tue Nov 03 2015 01:24:27 -0800: IKE SA negotiation successfully completed (4 times)
  Tue Nov 03 2015 01:23:38 -0800: User cleared IPSec SA from CLI (1 times)
  Tue Nov 03 2015 01:21:32 -0800: IPSec SA negotiation successfully completed (1 times)
  Tue Nov 03 2015 01:21:31 -0800: IPSec SA delete payload received from peer, corresponding
IPSec SAs cleared (1 times)
  Tue Nov 03 2015 01:21:27 -0800: IPSec SA negotiation successfully completed (1 times)
  Tue Nov 03 2015 01:21:13 -0800: Tunnel configuration changed. Corresponding IKE/IPSec SAs are
deleted (1 times)
  Tue Nov 03 2015 01:19:27 -0800: IPSec SA negotiation successfully completed (1 times)
  Tue Nov 03 2015 01:19:27 -0800: Tunnel is ready. Waiting for trigger event or peer to trigger
negotiation (1 times)
  Location: FPC 0, PIC 3, KMD-Instance 2
  Direction: inbound, SPI: 43de5d65, AUX-SPI: 0
  Hard lifetime: Expires in 1335 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 996 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes256-cbc (256 bits)
  Anti-replay service: counter-based enabled

  , Replay window size: 64
  Location: FPC 0, PIC 3, KMD-Instance 2
  Direction: outbound, SPI: 5b6e157c, AUX-SPI: 0
  Hard lifetime: Expires in 1335 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 996 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes256-cbc (256 bits)
```

```
  Anti-replay service: counter-based enabled

  , Replay window size: 64
```

show security ipsec security-associations detail (ADVPN Partner, Static Tunnel)

```
user@host> show security ipsec security-associations detail
ID: 67108872 Virtual-system: root, VPN Name: ZTH_SPOKE_VPN
  Local Gateway: 192.168.1.2, Remote Gateway: 192.168.1.1
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv2
  DF-bit: clear, Bind-interface: st0.1
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0x8608a29
  Tunnel events:
  Tue Nov 03 2015 01:24:26 -0800: IPSec SA negotiation successfully completed (1 times)
  Tue Nov 03 2015 01:24:26 -0800: IKE SA negotiation successfully completed (4 times)
  Tue Nov 03 2015 01:23:37 -0800: IPSec SA delete payload received from peer, corresponding
IPSec SAs cleared (1 times)
  Tue Nov 03 2015 01:21:31 -0800: IPSec SA negotiation successfully completed (1 times)
  Tue Nov 03 2015 01:21:31 -0800: Tunnel is ready. Waiting for trigger event or peer to trigger
negotiation (1 times)
  Tue Nov 03 2015 01:18:26 -0800: Key pair not found for configured local certificate.
Negotiation failed (1 times)
  Tue Nov 03 2015 01:18:13 -0800: CA certificate for configured local certificate not found.
Negotiation not initiated/successful (1 times)
  Direction: inbound, SPI: 5b6e157c, AUX-SPI: 0
  Hard lifetime: Expires in 941 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 556 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes256-cbc (256 bits)
  Anti-replay service: counter-based enabled, Replay window size: 64
  Direction: outbound, SPI: 43de5d65, AUX-SPI: 0
  Hard lifetime: Expires in 941 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 556 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes256-cbc (256 bits)
  Anti-replay service: counter-based enabled, Replay window size: 64
```

show security ipsec security-associations sa-type shortcut (ADVPN)

```
user@host> show security ipsec security-associations sa-type shortcut
Total active tunnels: 1
ID         Algorithm         SPI      Life:sec/kb  Mon lsys Port  Gateway
<268173318 ESP:aes256/sha256 6f164ee0 3580/ unlim - root 500 192.168.0.111
>268173318 ESP:aes256/sha256 e6f29cb0 3580/ unlim - root 500 192.168.0.111
```

show security ipsec security-associations sa-type shortcut detail (ADVPN)

```
user@host> show security ipsec security-associations sa-type shortcut detail
node0:
--------------------------------------------------------------------------

ID: 67108874 Virtual-system: root, VPN Name: ZTH_SPOKE_VPN
  Local Gateway: 192.168.1.2, Remote Gateway: 192.168.1.2
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Auto Discovery VPN:
    Type: Shortcut, Shortcut Role: Initiator
  Version: IKEv2
  DF-bit: clear, Bind-interface: st0.1
  Port: 4500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0x40608a29
  Tunnel events:
    Tue Nov 03 2015 01:47:26 -0800: IPSec SA negotiation successfully completed (1 times)
    Tue Nov 03 2015 01:47:26 -0800: Tunnel is ready. Waiting for trigger event or peer to
trigger negotiation (1 times)
    Tue Nov 03 2015 01:47:26 -0800: IKE SA negotiation successfully completed (1 times)
  Direction: inbound, SPI: b7a5518, AUX-SPI: 0
    Hard lifetime: Expires in 1766 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1381 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes256-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
  Direction: outbound, SPI: b7e0268, AUX-SPI: 0
    Hard lifetime: Expires in 1766 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1381 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
```

```
    Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes256-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
```

show security ipsec security-associations family inet detail

```
user@host> show security ipsec security-associations family inet detail
ID: 131073 Virtual-system: root, VPN Name: ike-vpn
  Local Gateway: 192.168.1.1, Remote Gateway: 192.168.1.2
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv1
  DF-bit: clear
  , Copy-Outer-DSCP Enabled
  Bind-interface: st0.99

  Port: 500, Nego#: 116, Fail#: 0, Def-Del#: 0 Flag: 0x600a29
  Tunnel events:
  Fri Oct 30 2015 15:47:21 -0700: IPSec SA rekey successfully completed (115 times)
  Fri Oct 30 2015 11:38:35 -0700: IKE SA negotiation successfully completed (12 times)
  Mon Oct 26 2015 16:41:07 -0700: IPSec SA negotiation successfully completed (1 times)
  Mon Oct 26 2015 16:40:56 -0700: Tunnel is ready. Waiting for trigger event or peer to trigger
negotiation (1 times)
  Mon Oct 26 2015 16:40:56 -0700: External interface's address received. Information updated (1
times)
  Location: FPC 0, PIC 1, KMD-Instance 1
  Direction: inbound, SPI: 81b9fc17, AUX-SPI: 0
  Hard lifetime: Expires in 1713 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 1090 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes256-cbc (256 bits)
  Anti-replay service: counter-based enabled

  , Replay window size: 64
  Location: FPC 0, PIC 1, KMD-Instance 1
  Direction: outbound, SPI: 727f629d, AUX-SPI: 0
  Hard lifetime: Expires in 1713 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 1090 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes256-cbc (256 bits)
```

```
  Anti-replay service: counter-based enabled

  , Replay window size: 64
```

show security ipsec security-associations detail (SRX4600)

```
user@host> show security ipsec security-associations detail
ID: 131073 Virtual-system: root, VPN Name: ike-vpn
  Local Gateway: 10.62.1.3, Remote Gateway: 10.62.1.2
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv2
  DF-bit: clear, Bind-interface: st0.0
  Port: 500, Nego#: 25, Fail#: 0, Def-Del#: 0 Flag: 0x600a29
  Tunnel events:
    Fri Jan 12 2007 07:50:10 -0800: IPSec SA rekey successfully completed (23 times)
  Location: FPC 0, PIC 0, KMD-Instance 0
  Anchorship: Thread 6
  Direction: inbound, SPI: 812c9c01, AUX-SPI: 0
    Hard lifetime: Expires in 2224 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1598 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes256-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
  Location: FPC 0, PIC 0, KMD-Instance 0
  Anchorship: Thread 7
  Direction: outbound, SPI: c4de0972, AUX-SPI: 0
    Hard lifetime: Expires in 2224 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1598 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha-256, Encryption: aes256-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
```

show security ipsec security-associations detail (SRX5400, SRX5600, SRX5800)

A new output field `IKE SA Index` corresponding to every IPsec SA within a tunnel is displayed under each IPsec SA information.

```
user@host> show security ipsec security-associations detail
ID: 500005 Virtual-system: root, VPN Name: 85BX5-OAM
  Local Gateway: 10.217.0.4, Remote Gateway: 10.200.254.118
  Traffic Selector Name: TS_DEFAULT
  Local Identity: ipv4(0.0.0.0-255.255.255.255)
  Remote Identity: ipv4(10.181.235.224-10.181.235.224)
  Version: IKEv2
  PFS group: N/A
  DF-bit: clear, Copy-Outer-DSCP Disabled, Bind-interface: st0.0, Policy-name: MACRO-IPSEC-POL
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0
  Multi-sa, Configured SAs# 0, Negotiated SAs#: 0
  Location: FPC 7, PIC 1, KMD-Instance 0
  Anchorship: Thread 15
  Distribution-Profile: default-profile
  Direction: inbound, SPI: 0xe2eb3838, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 644 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 159 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: aes128-gcm, Encryption: aes-gcm (128 bits)
    Anti-replay service: disabled
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-responder-only
    IKE SA Index: 22
  Direction: outbound, SPI: 0x4f7c3101, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 644 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 159 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: aes128-gcm, Encryption: aes-gcm (128 bits)
    Anti-replay service: disabled
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-responder-only
    IKE SA Index: 22
  Direction: inbound, SPI: 0x30b6d66f, AUX-SPI: 0
```

```
                              , VPN Monitoring: -
    Hard lifetime: Expires in 1771 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1391 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: aes128-gcm, Encryption: aes-gcm (128 bits)
    Anti-replay service: disabled
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-responder-only
          IKE SA Index: 40
Direction: outbound, SPI: 0xd2db4108, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 1771 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1391 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: aes128-gcm, Encryption: aes-gcm (128 bits)
    Anti-replay service: disabled
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-responder-only
          IKE SA Index: 40
```

show security ipsec security-associations ha-link-encryption (SRX5400, SRX5600, SRX5800)

Starting in Junos OS Release 20.4R1, when you configure the high availability (HA) feature, you can use this show command to view only interchassis link tunnel details.

```
user@host> show security ipsec security-associations ha-link-encryption
  Total active tunnels: 1     Total IPsec sas: 91
  ID      Algorithm       SPI      Life:sec/kb  Mon lsys Port  Gateway
  <495001 ESP:aes-gcm-256/aes256-gcm 0x0047658d 298/ unlim - root 500 10.23.0.2
  >495001 ESP:aes-gcm-256/aes256-gcm 0x0046c5cd 298/ unlim - root 500 10.23.0.2
  <495001 ESP:aes-gcm-256/aes256-gcm 0x0447658d 298/ unlim - root 500 10.23.0.2
  >495001 ESP:aes-gcm-256/aes256-gcm 0x0446c5cd 298/ unlim - root 500 10.23.0.2
  <495001 ESP:aes-gcm-256/aes256-gcm 0x0847658d 298/ unlim - root 500 10.23.0.2
  >495001 ESP:aes-gcm-256/aes256-gcm 0x0846c5cd 298/ unlim - root 500 10.23.0.2
  <495001 ESP:aes-gcm-256/aes256-gcm 0x0c47658d 298/ unlim - root 500 10.23.0.2
  >495001 ESP:aes-gcm-256/aes256-gcm 0x0c46c5cd 298/ unlim - root 500 10.23.0.2
  <495001 ESP:aes-gcm-256/aes256-gcm 0x1047658d 298/ unlim - root 500 10.23.0.2
  >495001 ESP:aes-gcm-256/aes256-gcm 0x1046c5cd 298/ unlim - root 500 10.23.0.2
```

```
  <495001 ESP:aes-gcm-256/aes256-gcm 0x1447658d 298/ unlim - root 500 10.23.0.2
  >495001 ESP:aes-gcm-256/aes256-gcm 0x1446c5cd 298/ unlim - root 500 10.23.0.2
  <495001 ESP:aes-gcm-256/aes256-gcm 0x1847658d 298/ unlim - root 500 10.23.0.2
  >495001 ESP:aes-gcm-256/aes256-gcm 0x1846c5cd 298/ unlim - root 500 10.23.0.2
  <495001 ESP:aes-gcm-256/aes256-gcm 0x1c47658d 298/ unlim - root 500 10.23.0.2
  >495001 ESP:aes-gcm-256/aes256-gcm 0x1c46c5cd 298/ unlim - root 500 10.23.0.2
  <495001 ESP:aes-gcm-256/aes256-gcm 0x2047658d 298/ unlim - root 500 10.23.0.2
  >495001 ESP:aes-gcm-256/aes256-gcm 0x2046c5cd 298/ unlim - root 500 10.23.0.2
  <495001 ESP:aes-gcm-256/aes256-gcm 0x2447658d 298/ unlim - root 500 10.23.0.2
  >495001 ESP:aes-gcm-256/aes256-gcm 0x2446c5cd 298/ unlim - root 500 10.23.0.2
...
```

show security ipsec sa detail ha-link-encryption (SRX5400, SRX5600, SRX5800)

Starting in Junos OS Release 20.4R1, when you configure the high availability (HA) feature, you can use this show command to view only interchassis link tunnel details. It displays the multi SAs created for interchassis link encryption tunnel.

```
user@host> show security ipsec sa detail ha-link-encryption
ID: 495001 Virtual-system: root, VPN Name: L3HA_IPSEC_VPN
  Local Gateway: 10.23.0.1, Remote Gateway: 10.23.0.2
  Traffic Selector Name: __L3HA_IPSEC_VPN__multi_node__
  Local Identity: ipv4(180.100.1.1-180.100.1.1)
  Remote Identity: ipv4(180.100.1.2-180.100.1.2)
  Version: IKEv2
  PFS group: DH-Group-24
  DF-bit: clear, Copy-Outer-DSCP Disabled, Bind-interface: st0.16000, Policy-name: L3HA_IPSEC_POL
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0
  Multi-sa, Configured SAs# 0, Negotiated SAs#: 0
  HA Link Encryption Mode: Multi-Node
  Location: FPC -, PIC -, KMD-Instance -
  Anchorship: Thread -
  Distribution-Profile: default-profile
  Direction: inbound, SPI: 0x00439cf8, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 294 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 219 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: aes256-gcm, Encryption: aes-gcm (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
```

```
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-immediately
    Location: FPC 1, PIC 0, KMD-Instance 0
    Anchorship: Thread 15
    IKE SA Index: 4294966297
  Direction: outbound, SPI: 0x004cfceb, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 294 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 219 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: aes256-gcm, Encryption: aes-gcm (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-immediately
    Location: FPC 1, PIC 0, KMD-Instance 0
    Anchorship: Thread 15
    IKE SA Index: 4294966297
  Direction: inbound, SPI: 0x04439cf8, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 294 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 219 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: aes256-gcm, Encryption: aes-gcm (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-immediately
    Location: FPC 1, PIC 0, KMD-Instance 0
    Anchorship: Thread 16
    IKE SA Index: 4294966297
  Direction: outbound, SPI: 0x044cfceb, AUX-SPI: 0
                              , VPN Monitoring: -

...
```

In Junos OS Release 22.3R1 and later, when you configure the Chassis Cluster HA control link encryption feature, you can execute the `show security ike sa ha-link-encryption detail`, `show security ipsec sa ha-link-encryption detail`, and `show security ipsec sa ha-link-encryption` commands to view the Chassis cluster control link encryption tunnel details.

show security ike sa ha-link-encryption detail

```
user@host> show security ike sa ha-link-encryption detail
IKE peer 10.2.0.1, Index 4294966274, Gateway Name: IKE_GW_HA_0
  Role: Initiator, State: UP
  Initiator cookie: ae5bcb5540d388a1, Responder cookie: 28bbae629ceb727f
  Exchange type: IKEv2, Authentication method: Pre-shared-keys
  Local gateway interface: em0
  Routing instance: __juniper_private1__
  Local: 10.7.0.2:500, Remote: 10.2.0.1:500
  Lifetime: Expires in 24856 seconds
  Reauth Lifetime: Disabled
  IKE Fragmentation: Enabled, Size: 576
  Remote Access Client Info: Unknown Client
  Peer ike-id: 10.2.0.1
  AAA assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-2
  Traffic statistics:
   Input  bytes  :               200644
   Output bytes  :               200644
   Input  packets:                 2635
   Output packets:                 2635
   Input  fragmented packets:       0
   Output fragmented packets:       0
  IPSec security associations: 6 created, 3 deleted
  Phase 2 negotiations in progress: 1
  IPSec Tunnel IDs: 495002
    Negotiation type: Quick mode, Role: Initiator, Message ID: 0
    Local: 10.7.0.2:500, Remote: 10.2.0.1:500
    Local identity: 10.7.0.2
    Remote identity: 10.2.0.1
    Flags: IKE SA is created
 IPsec SA Rekey CREATE_CHILD_SA exchange stats:
   Initiator stats:                                Responder stats:
    Request Out             : 1                    Request In             : 1
    Response In             : 1                    Response Out           : 1
    No Proposal Chosen In   : 0                    No Proposal Chosen Out : 0
    Invalid KE In           : 0                    Invalid KE Out         : 0
```

```
    TS Unacceptable In      : 0                     TS Unacceptable Out    : 0
    Res DH Compute Key Fail : 0                     Res DH Compute Key Fail: 0
    Res Verify SA Fail      : 0
    Res Verify DH Group Fail: 0
    Res Verify TS Fail      : 0
```

show security ipsec sa ha-link-encryption detail

```
user@host> show security ipsec sa ha-link-encryption detail
ID: 495002 Virtual-system: root, VPN Name: IPSEC_VPN_HA_0
  Local Gateway: 10.7.0.2, Remote Gateway: 10.2.0.1
  Traffic Selector Name: __IPSEC_VPN_HA_0__l2_chassis_clu
  Local Identity: ipv4(10.7.0.2-10.7.0.2)
  Remote Identity: ipv4(10.2.0.1-10.2.0.1)
  TS Type: traffic-selector
  Version: IKEv2
  PFS group: DH-group-24
  DF-bit: clear, Copy-Outer-DSCP Disabled, Bind-interface: st0.16000, Tunnel MTU: 0, Policy-
name: IPSEC_POL_HA_0
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0
  Multi-sa, Configured SAs# 0, Negotiated SAs#: 0
  HA Link Encryption Mode: L2 Chassis Cluster
  Location: FPC -, PIC -, KMD-Instance -
  Anchorship: Thread -
  Distribution-Profile: default-profile
  Direction: inbound, SPI: 0x35fae26b, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 3435 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2818 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-immediately
    IKE SA Index: 4294966274
  Direction: outbound, SPI: 0x0a2b9927, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 3435 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2818 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
```

```
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-immediately
    IKE SA Index: 4294966274
```

show security ipsec sa ha-link-encryption

```
user@host> show security ipsec sa ha-link-encryption
Total active tunnels: 1     Total IPsec sas: 1
  ID      Algorithm       SPI      Life:sec/kb  Mon lsys Port  Gateway
  <495002 ESP:aes-cbc-256/sha1 0x35fae26b 3484/ unlim - root 500 10.2.0.1
  >495002 ESP:aes-cbc-256/sha1 0x0a2b9927 3484/ unlim - root 500 10.2.0.1
```

show security ipsec security-associations detail (SRX Series devices and MX Series Routers)

In Junos OS Release 20.4R2, 21.1R1, and later, you can execute the `show security ipsec security-associations detail` command to view the traffic selector type for a VPN.

```
user@host> show security ipsec security-associations detail
ID: 500024 Virtual-system: root, VPN Name: S2S_VPN2
  Local Gateway: 10.7.0.2, Remote Gateway: 10.2.0.1
  Traffic Selector Name: ts1
  Local Identity: ipv4(10.20.20.0-10.20.20.255)
  Remote Identity: ipv4(10.10.10.0-10.10.10.255)
  TS Type: traffic-selector
  Version: IKEv2
  PFS group: DH-group-14
  DF-bit: clear, Copy-Outer-DSCP Disabled, Bind-interface: st0.2, Policy-name: IPSEC_POL
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0
  Multi-sa, Configured SAs# 0, Negotiated SAs#: 0
  Tunnel events:
    Tue Jan 19 2021 04:43:49: IPsec SA negotiation succeeds (1 times)
  Location: FPC 0, PIC 0, KMD-Instance 0
  Anchorship: Thread 1
  Distribution-Profile: default-profile
  Direction: inbound, SPI: 0xf8642fae, AUX-SPI: 0
```

```
                              , VPN Monitoring: -
    Hard lifetime: Expires in 1798 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1397 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha256-128, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-immediately
    IKE SA Index: 17
  Direction: outbound, SPI: 0xb2a26969, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 1798 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1397 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha256-128, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-immediately
    IKE SA Index: 17
ID: 500025 Virtual-system: root, VPN Name: S2S_VPN1
  Local Gateway: 10.7.0.1, Remote Gateway: 10.2.0.1
  Local Identity: ipv4(0.0.0.0-255.255.255.255)
  Remote Identity: ipv4(0.0.0.0-255.255.255.255)
  TS Type: proxy-id
  Version: IKEv2
  PFS group: DH-group-14
  DF-bit: clear, Copy-Outer-DSCP Disabled, Bind-interface: st0.1, Policy-name: IPSEC_POL
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0
  Multi-sa, Configured SAs# 0, Negotiated SAs#: 0
  Tunnel events:
    Tue Jan 19 2021 04:44:41: IPsec SA negotiation succeeds (1 times)
  Location: FPC 0, PIC 0, KMD-Instance 0
  Anchorship: Thread 1
  Distribution-Profile: default-profile
  Direction: inbound, SPI: 0xe293762a, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 1755 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1339 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha256-128, Encryption: aes-cbc (256 bits)
```

```
    Anti-replay service: counter-based enabled, Replay window size: 64
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-immediately
    IKE SA Index: 18
  Direction: outbound, SPI: 0x7aef9d7f, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 1755 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1339 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha256-128, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
    Extended-Sequence-Number: Disabled
    tunnel-establishment: establish-tunnels-immediately
    IKE SA Index: 18
```

show security ipsec security-associations detail (SRX5400, SRX5600, SRX5800)

Starting in Junos OS Release 21.1R1, you can view the traffic selector details, that includes, local identity, remote identity, protocol, source-port range, destination port range for multiple terms defined for an IPsec SA.

In the earlier Junos Releases, traffic selection for a particular SA is performed using existing IP range defined using IP address or netmask. From Junos OS Release 21.1R1 onwards, additionally traffic is selected through protocol specified using *protocol_name*. And also, low and high port range specified for source and destination port numbers.

```
user@host> show security ipsec security-associations detail

ID: 500075 Virtual-system: root, VPN Name: pkn-r0-r1-ipsec-vpn-1
Local Gateway: 10.1.1.1, Remote Gateway: 10.1.1.2

Traffic Selector Name: ts1

   Local Identity:
   Protocol         Port                IP
   17/UDP         100-200             198.51.100.0-198.51.100.255
   6/TCP          250-300             198.51.100.0-198.51.100.255
   Remote Identity:
   Protocol         Port                IP
   17/UDP         150-200             10.80.0.1-10.80.0.1
   6/TCP          250-300             10.80.1.1-10.80.1.1
```

```
Version: IKEv2
DF-bit: clear, Copy-Outer-DSCP Disabled, Bind-interface: st0.0, Policy-name: pkn-r0-r1-ipsec-
policy
Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0
Multi-sa, Configured SAs# 0, Negotiated SAs#: 0
Location: FPC 0, PIC 0, KMD-Instance 0
Anchorship: Thread 1
Distribution-Profile: default-profile
Direction: inbound, SPI: ……
Direction: outbound, SPI: ……
```

show security ipsec security-associations srg-id

```
user@host> show security ipsec security-associations srg-id 1

Total active tunnels: 1     Total IPsec sas: 2
  ID      Algorithm       SPI      Life:sec/kb  Mon lsys Port  Gateway
  <17277217 ESP:aes-cbc-256/sha256 0xc7faee3e 1440/ unlim - root 500 10.112.0.1
  >17277217 ESP:aes-cbc-256/sha256 0x7921d472 1440/ unlim - root 500 10.112.0.1
  <17277217 ESP:aes-cbc-256/sha256 0xf1a01dd4 1498/ unlim - root 500 10.112.0.1
  >17277217 ESP:aes-cbc-256/sha256 0xa0b77273 1498/ unlim - root 500 10.112.0.1
```

Release Information

Command introduced in Junos OS Release 8.5. Support for the `family` option added in Junos OS Release 11.1.

Support for the `vpn-name` option added in Junos OS Release 11.4R3. Support for the `traffic-selector` option and traffic selector field added in Junos OS Release 12.1X46-D10.

Support for Auto Discovery VPN (ADVPN) added in Junos OS Release 12.3X48-D10.

Support for IPsec datapath verification added in Junos OS Release 15.1X49-D70.

Support for thread anchorship added in Junos OS Release 17.4R1.

Starting in Junos OS Release 18.2R2 the `show security ipsec security-assocations detail` command output will include thread anchorship information for the security associations (SAs).

Starting in Junos OS Release 19.4R1, we have deprecated the CLI option `fc-name` (COS Forward Class name) in the new **iked** process that displays the security associations (SAs) under show command `show security ipsec sa`.

Support for the `ha-link-encryption` option added in Junos OS Release 20.4R1.

Support for the `srg-id` option added in Junos OS Release 22.4R1.

Support for `passive-mode-tunneling` on MX-SPC3 is introduced in Junos OS Release 23.1R1.

RELATED DOCUMENTATION

Example: Configuring a Route-Based VPN Tunnel in a User Logical Systems

show security log report

IN THIS SECTION

- Syntax | **1209**
- Description | **1210**
- Options | **1210**
- Required Privilege Level | **1210**
- Output Fields | **1210**
- Sample Output | **1210**
- Release Information | **1211**

Syntax

```
in-detail
in-interval
summary
top
```

Description

Displays the security log report settings.

On-box reporting offers a comprehensive reporting facility where your security management team can spot a security event when it occurs. Immediately access and review pertinent details about the event, and quickly decide appropriate remedial action.

Options

in-detail	Displays the detail log content
in-interval	Displays the count in intervals
summary	Displays the summary information
top	Displays the top number to be calculated

Required Privilege Level

view

Output Fields

Sample Output

show security log report summary session-all tenant TSYS1

```
user@host> show security log report summary session-all tenant TSYS1
total-count: 4
```

Show security log report summary session-all root-logical-system

```
user@host> Show security log report summary session-all root-logical-system
total-count: 4
```

Release Information

Command introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

exclude

clear security log

show security match-policies

IN THIS SECTION

- Syntax | **1212**
- Description | **1212**
- Options | **1213**
- Required Privilege Level | **1214**
- Output Fields | **1214**
- Sample Output | **1217**
- Release Information | **1220**

Syntax

```
show security match-policies
destination-ip <ip-address>
destination-port < port-number>
destination-vrf <destination-vrf>
from-zone <zone-name>
global
logical-system <logical-system-name>
protocol <protocol-name | protocol-number>
result-count <number>
root-logical-system
source-end-user-profile <device-identity-profile-name>
source-identity <role-name>
source-ip <ip-address>
source-port <port-number>
tenant <tenant-name>
to-zone <zone-name>
url category<url-category>
```

Description

The `show security match-policies` command allows you to troubleshoot traffic problems using the match criteria: source port, destination port, source IP address, destination IP address, and protocol. For example, if your traffic is not passing because either an appropriate policy is not configured or the match criteria is incorrect, then the `show security match-policies` command allows you to work offline and identify where the problem actually exists. It uses the search engine to identify the problem and thus enables you to use the appropriate match policy for the traffic.

The `result-count` option specifies how many policies to display. The first enabled policy in the list is the policy that is applied to all matching traffic. Other policies below it are “shadowed” by the first and are never encountered by matching traffic.

NOTE: The `show security match-policies` command is applicable only to security policies; IDP policies are not supported.

Options

- destination-ip *destination-ip*—Displays the destination IP address of the traffic.
- destination-port *destination-port*–Displays the destination port number of the traffic. Range is 1 through 65,535.
- destination-vrf *destination-vrf*—(Optional) Displays the destination VRF information.
- from-zone *zone-name*—Displays the name or ID of the source zone of the traffic. This is optional for global policies.
- global—Displays information about global policies.
- logical-system—Displays the logical system name.
- protocol *protocol-name* | *protocol-number*–Displays the protocol name or numeric value of the traffic.
  - ah or 51
  - egp or 8
  - esp or 50
  - gre or 47
  - icmp or 1
  - igmp or 2
  - igp or 9
  - ipip or 94
  - ipv6 or 41
  - ospf or 89
  - pgm or 113
  - pim or 103
  - rdp or 27
  - rsvp or 46
  - sctp or 132
  - tcp or 6

- udp or 17
- vrrp or 112

- result-count *number*—(Optional) Displays the number of policy matches. Valid range is from 1 through 16. The default value is 1.
- root-logical-system—Displays root logical system as default.
- source-end-user-profile *device-identity-profile-name*—(Optional) Displays the device identity profile that specifies characteristics that can apply to one or more devices.
- source-identity *role-name*—(Optional) Displays the source identity of the traffic determined by the user role.
- source-ip *source-ip*—Displays the source IP address of the traffic.
- source-port *source-port*—Displays the source port number of the traffic. Range is 1 through 65,535.
- tenant—Displays the name of the tenant system.
- to-zone *zone-name*—Displays the name or ID of the destination zone of the traffic. This is optional for global policies.

Required Privilege Level

view

Output Fields

Table 73 on page 1214 lists the output fields for the show security match-policies command. Output fields are listed in the approximate order in which they appear.

Table 73: show security match-policies Output Fields

Field Name	Field Description
Policy	Name of the applicable policy.

Table 73: show security match-policies Output Fields *(Continued)*

Field Name	Field Description
`Action` or `Action-type`	The action to be taken for traffic that matches the policy's match criteria. Actions include the following: • `permit` • `firewall-authentication` • `tunnel ipsec-vpn` *`vpn-name`* • `pair-policy` *`pair-policy-name`* • `deny` • `reject`
`State`	Status of the policy: • `enabled`: The policy can be used in the policy lookup process, which determines access rights for a packet and the action taken in regard to it. • `disabled`: The policy cannot be used in the policy lookup process, and therefore it is not available for access control.
`Index`	An internal number associated with the policy.
`Sequence number`	Number of the policy within a given context. For example, three policies that are applicable in a from-zoneA-to-zoneB context might be ordered with sequence numbers 1, 2, and 3. Also, in a from-zoneC-to-zoneD context, four policies might have sequence numbers 1, 2, 3, and 4.
`From zone`	Name of the source zone.
`To zone`	Name of the destination zone.
`Source addresses`	The names and corresponding IP addresses of the source addresses for a policy. Address sets are resolved to their individual address name-IP address pairs.

Table 73: show security match-policies Output Fields *(Continued)*

Field Name	Field Description
Destination addresses	The names and corresponding IP addresses of the destination addresses (or address sets) for a policy as entered in the destination zone's address book. A packet's destination address must match one of these addresses for the policy to apply to it.
Application	Name of a preconfigured or custom application, or any if no application is specified.
IP protocol	Numeric value for the IP protocol used by the application, such as 6 for TCP or 1 for ICMP.
ALG	If an ALG is associated with the session, the name of the ALG. Otherwise, 0.
Inactivity timeout	Elapsed time without activity after which the application is terminated.
Source-port range	Range of matching source ports defined in the policy.
Destination-port range	Range of matching destination ports defined in the policy.
Source identities	One or more user roles defined in the matching policy.
global	Display information about global policies.
device-identity-profile-name	Device identity profile that specifies characteristics that can apply to one or more devices.

Sample Output

Example 1: show security match-policies

```
user@host> show security match-policies from-zone z1 to-zone z2 source-ip 10.10.10.1 destination-
ip 192.0.2.1 source-port 1 destination-port 21 protocol tcp
Policy: p1, action-type: permit, State: enabled, Index: 4
  Sequence number: 1
  From zone: z1, To zone: z2
  Source addresses:
    a2: 198.51.100.0/24
    a3: 10.10.10.1/32
  Destination addresses:
    d2: 203.0.113.0/24
    d3: 192.0.2.1/32
  Application: junos-ftp
    IP protocol: tcp, ALG: ftp, Inactivity timeout: 1800
            Source port range: [0-0]
            Destination port range: [21-21]
```

Example 2: show security match policies ... result-count

```
user@host> show security match-policies from-zone zone-A to-zone zone-B source-ip 10.10.10.1
destination-ip 192.0.2.5 source_port 1004 destination_port 80 protocol tcp result_count 5
Policy: p1, action-type: permit, State: enabled, Index: 4
  Sequence number: 1
  From zone: zone-A, To zone: zone-B
  Source addresses:
    sa1: 10.10.0.0/16
  Destination addresses:
    da5: 192.0.2.0/24
  Application: any
    IP protocol: 1, ALG: 0, Inactivity timeout: 0
            Source port range: [1000-1030]
          Destination port range: [80-80]

Policy: p15, action-type: deny, State: enabled, Index: 18
  Sequence number: 15
  From zone: zone-A, To zone: zone-B
  Source addresses:
```

```
    sa11: 10.10.10.1/32
  Destination addresses:
    da15: 192.0.2.5/32
  Application: any
    IP protocol: 1, ALG: 0, Inactivity timeout: 0
            Source port range: [1000-1030]
          Destination port range: [80-80]
```

Example 3: show security match policies ... source-identity

```
user@host> show security match-policies from-zone untrust to-zone trust  source-ip 10.10.10.1
destination-ip 192.0.2.1 destination_port 21 protocol 6 source-port 1234 source-identity role1
Policy: p1, action-type: permit, State: enabled, Index: 40
      Policy Type: Configured
      Sequence number: 1
      From zone: untrust, To zone: trust
      Source addresses:
        a1: 10.0.0.0/8
      Destination addresses:
        d1: 192.0.2.0/24
      Application: junos-ftp
        IP protocol: tcp, ALG: ftp, Inactivity timeout: 1800
          Source port range: [0-0]
          Destination port range: [21-21]
      Source identities: role1
      Per policy TCP Options: SYN check: No, SEQ check: No
```

Example 4: show security match policies ... global

```
user@host> show security match-policies global source-ip 10.10.10.1 destination-ip 192.0.2.5
source_port 1004 destination_port 80 protocol tcp result_count 5
Policy: gp1, action-type: permit, State: enabled, Index: 6, Scope Policy: 0
  Policy Type: Configured, global
  Sequence number: 1
  From zones:
    Any
  To zones:
    Any
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
```

```
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No, Window scale: No
```

show security match-policies tenant TN1 from-zone trust to-zone untrust source-ip 10.10.10.1 destination-ip 192.0.2.1 source-port 1 destination-port 21 protocol tcp

```
user@host> show security match-policies tenant TN1 from-zone trust to-zone untrust source-ip
10.10.10.1 destination-ip 192.0.2.1 source-port 1 destination-port 21 protocol tcp

Policy: p1, action-type: permit, State: enabled, Index: 4
Sequence number: 1
From zone: trust, To zone: untrust
Source addresses:
a2: 198.51.100.0/24
a3: 10.10.10.1/32
Destination addresses:
d2: 203.0.113.0/24
d3: 192.0.2.1/32
Application: junos-ftp
IP protocol: tcp, ALG: ftp, Inactivity timeout: 1800
Source port range: [0-0]
Destination port range: [21-21]
```

show security match-policies from-zone client to-zone svr source-ip 10.1.1.1 source-port 88 destination-ip 10.2.2.2 destination-port 80 protocol tcp url-category Enhanced_Games

```
user@host> show security match-policies from-zone client to-zone svr source-ip 10.1.1.1 source-
port 88 destination-ip 10.2.2.2 destination-port 80 protocol tcp url-category Enhanced_Games

Policy: p1, action-type: permit, State: enabled, Index: 7
0
  Policy Type: Configured
```

```
Sequence number: 1
From zone: client, To zone: server
Source vrf group:
  any
Destination vrf group:
  any
Source addresses:
  any-ipv4(global): 0.0.0.0/0
  any-ipv6(global): ::/0
Destination addresses:
  any-ipv4(global): 0.0.0.0/0
  any-ipv6(global): ::/0
Application: any
  IP protocol: 0, ALG: 0, Inactivity timeout: 0
    Source port range: [0-0]
    Destination ports: [0-0]
Url-category:
  Enhanced_Sex: 234881056
  Enhanced_Games: 234881037
Per policy TCP Options: SYN check: No, SEQ check: No, Window scale: No
Intrusion Detection and Prevention: disabled
Unified Access Control: disabled
Unified Threat Management: enabled
```

Release Information

Command introduced in Junos OS Release 10.3.

Command updated in Junos OS Release 10.4.

Command updated in Junos OS Release 12.1.

Command updated to include optional from-zone and to-zone global match options in Junos OS Release 12.1X47-D10.

The `tenant` option is introduced in Junos OS Release 18.3R1.

The `url category` option is introduced in Junos OS Release 20.2R1.

RELATED DOCUMENTATION

- *clear security policies statistics*
- *Security Policies Overview*
- *Understanding Security Policy Rules*
- *Understanding Security Policy Elements*

show security nat destination rule

IN THIS SECTION

- Syntax | **1221**
- Description | **1222**
- Options | **1222**
- Required Privilege Level | **1222**
- Output Fields | **1222**
- Sample Output | **1224**
- Sample Output | **1225**
- Release Information | **1226**

Syntax

```
show security nat destination rule
rule-name
all
logical-system (logical-system-name)
root-logical-system
tenant (tenant-name)
```

Description

Display information about the specified destination Network Address Translation (NAT) rule. Destination NAT rules are processed after static NAT rules but before source NAT rules.

Options

rule-name	Display information about the specified destination NAT rule.
all	Display information about all the destination NAT rules.
logical-system	Display information about the destination NAT rules for a specified logical system. Specify `all` to display information for all logical systems.
root-logical-system	Display information about the destination NAT rules for the primary (root) logical system.
tenant	Display information about the destination NAT rules for a specified tenant system. Specify `all` to display information for all tenant systems.

Required Privilege Level

view

Output Fields

Table 74 on page 1222 lists the output fields for the **show security nat destination rule** command. Output fields are listed in the approximate order in which they appear.

Table 74: show security nat destination rule Output Fields

Field Name	Field Description
Total destination-nat rules	Number of destination NAT rules.

Table 74: show security nat destination rule Output Fields *(Continued)*

Field Name	Field Description
Total referenced IPv4/IPv6 ip-prefixes	Number of IP prefixes referenced in source, destination, and static NAT rules. This total includes the IP prefixes configured directly as address names and as address set names in the rule.
Destination NAT rule	Name of the destination NAT rule.
Description	Description of the destination NAT rule.
Rule-Id	Rule identification number.
Rule position	Position of the destination NAT rule.
From routing instance	Name of the routing instance from which the packets flow.
From interface	Name of the interface from which the packets flow.
From zone	Name of the zone from which the packets flow.
Source addresses	Name of the source addresses that match the rule. The default value is any.
Destination addresses	Name of the destination addresses that match the rule. The default value is any.
Action	The action taken when a packet matches the rule's tuples. Actions include the following: • **destination NAT pool**—Use user-defined destination NAT pool to perform destination NAT. • **off**—Do not perform destination NAT.
Destination ports	Destination ports number that match the rule. The default value is any.

Table 74: show security nat destination rule Output Fields *(Continued)*

Field Name	Field Description
Application	Indicates whether the application option is configured.
Translation hits	Number of translation hits.
Successful sessions	Number of successful session installations after the NAT rule is matched.
Failed sessions	Number of unsuccessful session installations after the NAT rule is matched.
Number of sessions	Number of sessions that reference the specified rule.

Sample Output

show security nat destination rule namelength-is-now-changed-upto-63-characterslength

```
user@host>show security nat destination rule namelength-is-now-changed-upto-63-characterslength

 Destination NAT rule: namelength-is-now-changed-upto-63-characterslength           Rule-set:
namelength-is-now-changed-upto-63-characterslength
  Description               : The destination rule dst2-rule is for the sales team
  Rule-Id                   : 1
  Rule position             : 1
  From routing instance     : ri1
                            : ri2
  Match
    Source addresses        : add1
                              add2
    Destination addresses   : add9
  Action                    : off

  Destination port          : 0
 Translation hits           : 68
   Successful sessions      : 25
```

```
    Failed sessions          : 43
  Number of sessions         : 2
```

Sample Output

show security nat destination rule all

```
user@host> show security nat destination rule all

Total destination-nat rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 2/0

Destination NAT rule: r4                  Rule-set: rs4
  Rule-Id                   : 2
  Rule position             : 2
  From zone                 : untrust
  Match
    Source addresses        : 192.0.2.0 - 192.0.2.255
    Destination addresses   : 198.51.100.0 - 198.51.100.255
    Application             : configured
  Action                    : off
  Translation hits          : 0
    Successful sessions     : 0
    Failed sessions         : 0
  Number of sessions        : 0
```

show security nat destination rule all tenant

```
user@host> show security nat destination rule all tenant tn1

Total destination-nat rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 2/0
Destination NAT rule: r1                    Rule-set: from_zone
  Rule-Id                   : 1
  Rule position             : 1
  From zone                 : untrust
```

```
Match
  Source addresses         : 192.0.2.0       - 192.0.2.255
  Destination addresses    : 203.0.113.202   - 203.0.113.202
Action                     : h1
Translation hits           : 0
  Successful sessions      : 0
  Failed sessions          : 0
Number of sessions         : 0
```

Release Information

Command introduced in Junos OS Release 9.2. The `Description` output field added in Junos OS Release 12.1.

Support for IPv6 logical systems and the `Successful sessions`, `Failed sessions` and `Number of sessions` output fields added in Junos OS Release 12.1X45-D10.

Output for multiple destination ports and the `application` option field added in Junos OS Release 12.1X47-D10.

The tenant option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

rule (Security Destination NAT)

show security nat destination summary

IN THIS SECTION

- Syntax | **1227**
- Description | **1227**
- Options | **1227**
- Required Privilege Level | **1228**

Output Fields | 1228
Sample Output | 1229
Release Information | 1230

Syntax

```
show security nat destination summary
logical-system (logical-system-name )
root-logical-system
tenant (tenant-name )
```

Description

Display a summary of Network Address Translation (NAT) destination pool information.

Options

none	Display summary information about the destination NAT pool.
logical-system	Display summary information about the destination NAT for a specified logical system. Specify `all` to display information for all logical systems.
root-logical-system	Display summary information about the destination NAT for the primary (root) logical system.
tenant	Display information about the destination NAT for a specified tenant system. Specify `all` to display information for all tenant systems.

Required Privilege Level

view

Output Fields

Table 75 on page 1228 lists the output fields for the **show security nat destination summary** command. Output fields are listed in the approximate order in which they appear.

Table 75: show security nat destination summary Output Fields

Field Name	Field Description
Total destination nat pool number	Number of destination NAT pools.
Pool name	Name of the destination address pool.
Address range	IP address or IP address range for the pool.
Routing Instance	Name of the routing instance.
Port	Port number.
Total	Number of IP addresses that are in use.
Available	Number of IP addresses that are free for use.
Total destination nat rule number	Number of destination NAT rules.
Total hit times	Number of times a translation in the translation table is used for all the destination NAT rules.

Table 75: show security nat destination summary Output Fields *(Continued)*

Field Name	Field Description
Total fail times	Number of times a translation in the translation table failed to translate for all the destination NAT rules.

Sample Output

show security nat destination summary

```
user@host> show security nat destination summary
Total pools: 1
Pool name           Address                          Routing        Port  Total
                    Range                            Instance             Address
dst_v4_pool-namelength-is-now-changed-upto-63-characters-length 192.0.2.2-192.0.2.2 0 1

Total rules: 1
Rule name : dstrule_namelen_is_now_changed_upto_63_characters_length
    Rule set : dst_ruleset_namelength_changed_upto_63_characters_length
    Action   : dst_v4_pool-namelength-is-now-changed-upto-63-characters-length
    From     : trust
```

show security nat destination summary tenant

```
user@host> show security nat destination summary tenant tn1

Total pools: 1
Pool name           Address                          Routing        Port  Total
                    Range                            Instance             Address
h1                  192.168.1.200  - 192.168.1.200                0     1

Total rules: 1
Rule name           Rule set       From                            Action
r1                  from_zone      untrust                         h1
```

Release Information

Command introduced in Junos OS Release 9.2.

Support for IPv6 logical systems added in Junos OS Release 12.1X45-D10.

The tenant option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

pool (Security Destination NAT)

rule (Security Destination NAT)

show security nat source rule

IN THIS SECTION

- Syntax | **1230**
- Description | **1231**
- Options | **1231**
- Required Privilege Level | **1231**
- Output Fields | **1231**
- Sample Output | **1234**
- Sample Output | **1234**
- Sample Output | **1236**
- Release Information | **1236**

Syntax

```
show security nat static rule
rule-name
```

```
all
logical-system (logical-system-name )
root-logical-system
tenant (tenant-name )
```

Description

Display information about the specified source Network Address Translation (NAT) rule.

Options

rule-name	Name of the rule.
all	Display information about all the source NAT rules.
logical-system	Display information about the source NAT rules for a specified logical system. Specify `all` to display information for all logical systems.
root-logical-system	Display information about the source NAT rules for the primary (root) logical system.
tenant	Display information about the source NAT rules for a specified tenant system. Specify `all` to display information for all tenant systems.

Required Privilege Level

view

Output Fields

Table 76 on page 1232 lists the output fields for the **show security nat source rule** command. Output fields are listed in the approximate order in which they appear

Table 76: show security nat source rule Output Fields

Field Name	Field Description
Source NAT rule	Name of the source NAT rule.
Total rules	Number of source NAT rules.
Total referenced IPv4/IPv6 ip-prefixes	Number of IP prefixes referenced in source, destination, and static NAT rules. This total includes the IP prefixes configured directly, as address names, and as address set names in the rule.
Description	Description of the source NAT rule.
Rule-Id	Rule identification number.
Rule position	Position of the source NAT rule.
From zone	Name of the zone from which the packets flow.
To zone	Name of the zone to which the packets flow.
From routing instance	Name of the routing instance from which the packets flow.
To routing instance	Name of the routing instance to which the packets flow.
From interface	Name of the interface from which the packets flow.
To interface	Name of the interface to which the packets flow.
Source addresses	Name of the source addresses that match the rule.
Source port	Source port numbers that match the rule.
Destination address	Name of the destination addresses that match the rule.

Table 76: show security nat source rule Output Fields *(Continued)*

Field Name	Field Description
Destination ports	Destination port numbers that match the rule.
Application	Indicates whether the application option is configured.
Action	The action taken in regard to a packet that matches the rule's tuples. Actions include the following: • **off**—Do not perform source NAT. • **source NAT pool**—Use user-defined source NAT pool to perform source NAT • **interface**—Use egress interface's IP address to perform source NAT.
Persistent NAT type	Persistent NAT type.
Persistent NAT mapping type	Persistent NAT mapping type.
Inactivity timeout	Inactivity timeout for persistent NAT binding.
Max session number	Maximum number of sessions.
Translation hits	Number of translation hits.
Successful sessions	Number of successful session installations after the NAT rule is matched.
Failed sessions	Number of unsuccessful session installations after the NAT rule is matched.
Number of sessions	Number of sessions that reference the specified rule.

Sample Output

show security nat source rule rule_namelength_is_now_changed_upto_63_characters_length

```
user@host> show security nat source rule rule_namelength_is_now_changed_upto_63_characters_length

source NAT rule: rule_namelength_is_now_changed_upto_63_characters_length
  Rule set                  : ruleset_namelength_now_changed_upto_63_characters_length
  Rule Id                   : 1
  Rule position             : 1
  From zone                 : trust
  To zone                   : untrust
  Match
    Source addresses        : 192.0.2.0 - 192.0.2.255
  Action                       : src_v4_pool-namelength-is-now-changed-upto-63-characters-length
    Persistent NAT type       : N/A
    Persistent NAT mapping type : address-port-mapping
    Inactivity timeout        : 0
    Max session number        : 0
  Translation hits          : 0
    Successful sessions     : 0
  Number of sessions        : 0
```

Sample Output

show security nat source rule all (SRX Series devices)

```
user@host> show security nat source rule all
Logical system: root
Total rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 3/0

source NAT rule: r2                  Rule-set: rs2
  Rule-Id                   : 2
  Rule position             : 1
  From zone                 : trust
  To zone                   : untrust
  Match
```

```
    Source addresses         : 192.0.2.0 - 192.0.2.255
    Destination addresses    : 203.0.113.0 - 203.0.113.255
                               198.51.100.0 - 198.51.100.255
    Application              : configured
  Action                       : off
    Persistent NAT type        : N/A
    Persistent NAT mapping type : address-port-mapping
    Inactivity timeout         : 0
    Max session number         : 0
  Translation hits           : 0
    Successful sessions      : 0
    Failed sessions          : 0
  Number of sessions         : 0
```

show services nat source rule all (MX-SPC3)

```
user@host> show services nat source rule all
Total rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 1/0
Interface: vms-4/0/0 , Service set: ss1
source NAT rule: NAT_RULE1--WEB-MX-SPC3NAT_WEB_MX_SPC3NAT_web-mx--spc3_NAT-RULE
  Rule set                   : NAT_RULE-SET1--WEB-MX-SPC3NAT_WEB_MX_SPC3NAT_web-mx--spc3_RSET
  Rule Id                    : 1
  Rule position              : 1
  Match
    Source addresses         : 0.0.0.0         - 255.255.255.255
    Application              : configured
  Action                       : NAT_POOL1--WEB-MX-SPC3NAT_WEB_MX_SPC3NAT_web-mx--spc3_NAT-POOL
    Syslog                     : enabled
    Mapping-type               : N/A
  Translation hits           : 0
    Successful sessions      : 0
    Failed sessions          : 0
  Number of sessions         : 0
```

Sample Output

show security nat source rule all tenant

```
user@host> show security nat source rule all tenant tn1
Total rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 2/0
source NAT rule: r1                     Rule-set: from_intf
  Rule-Id                   : 1
  Rule position             : 1
  From interface            : ge-0/0/0.0
  To interface              : ge-0/0/1.0
  Match
    Source addresses        : 192.168.1.0     - 192.168.1.255
    Destination addresses   : 203.0.113.200   - 203.0.113.200
  Action                       : pat
    Persistent NAT type        : N/A
    Persistent NAT mapping type : address-port-mapping
    Inactivity timeout         : 0
    Max session number         : 0
  Translation hits          : 0
    Successful sessions     : 0
    Failed sessions         : 0
  Number of sessions        : 0
```

Release Information

Command introduced in Junos OS Release 9.2. Support.

The **Description** output field added in Junos OS Release 12.1.

Support for IPv6 logical systems and the **Source port**, **Successful sessions**, **Failed sessions**, and **Number of sessions** output fields added in Junos OS Release 12.1X45-D10.

Output for multiple destination ports and the `application` output field added in Junos OS Release 12.1X47-D10.

The tenant option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

rule (Security Source NAT)

show security nat source summary

IN THIS SECTION

- Syntax | **1237**
- Description | **1237**
- Options | **1238**
- Required Privilege Level | **1238**
- Output Fields | **1238**
- Sample Output | **1239**
- show security nat source summary (SRX Series devices) | **1241**
- Release Information | **1244**

Syntax

```
show security nat source summary
logical-system (logical-system-name )
root-logical-system
tenant (tenant-name )
```

Description

Display a summary of Network Address Translation (NAT) source information.

Options

none	Display summary source NAT information.
logical-system	Display summary information about the source NAT for a specified logical system. Specify `all` to display information for all logical systems.
root-logical-system	Display summary information about the source NAT for the primary (root) logical system.
tenant	Display summary information about the source NAT for a specified tenant system. Specify `all` to display information for all tenant systems.

Required Privilege Level

view

Output Fields

Table 77 on page 1238 lists the output fields for the **show security nat source summary** command. Output fields are listed in the approximate order in which they appear.

Table 77: show security nat source summary Output Fields

Field Name	Field Description
Total source nat pool number	Number of source NAT pools.
Pool name	Name of the source address pool.
Address range	IP address or IP address range for the pool.
Routing Instance	Name of the routing instance.

Table 77: show security nat source summary Output Fields *(Continued)*

Field Name	Field Description
PAT	Whether Port Address Translation (PAT) is enabled (yes or no).
Total Address	Number of IP addresses that are in use. Starting in Junos OS Release 12.3X48-D55, and Junos OS Release 15.1X49-D90, and Junos OS Release 17.3R1, the total number of addresses that are in use for pools with IPv6 prefixes is shown as zero (0).
Total source nat rule number	Number of source NAT rules.
Total port number usage for port translation pool	Number of ports assigned to the pool.
Maximum port number for port translation pool	Maximum number of NAT or PAT transactions done at any given time.

Sample Output

show security nat source summary

```
user@host> show security nat source summary logical-system all


Logical system: root-logical-system
Total port number usage for port translation pool: 67108864
Maximum port number for port translation pool: 134217728

Logical system: lsys1
Total port number usage for port translation pool: 193536
Maximum port number for port translation pool: 134217728
Total pools: 2

Logical system: root-logical-system
```

```
Pool              Address                  Routing   PAT   Total
Name              Range                    Instance        Address
pool1             10.1.1.0-10.1.4.255-     default   yes   2048
                  10.1.5.0-10.1.8.255


Logical system: lsys1
Pool              Address                   Routing   PAT   Total
Name              Range                     Instance        Address
pool2             203.0.203.1-203.0.203.3    default   yes   3


Total rules: 1

Logical system: root-logical-system
Rule name         Rule set      From         To             Action
rule 1            ruleset1      ge-2/2/2.0   ge-2/2/3.0     pool1
rule 1                          ge-2/2/4.0   ge-2/2/5.0
```

show security nat source summary tenant

```
user@host> show security nat source summary tenant tn1

Total port number usage for port translation pool: 1548288
Maximum port number for port translation pool: 268435456
Total pools: 1
Pool                Address                 Routing             PAT  Total
Name                Range                   Instance                 Address
pat                 192.0.2.1-192.0.2.24    default             yes  24

Total rules: 1
Rule name         Rule set       From             To                  Action
r1                from_intf      ge-0/0/0.0       ge-0/0/1.0          pat
```

show security nat source summary (SRX Series devices)

```
user@host> show security nat source summary
Total port number usage for port translation pool: 4128768
Maximum port number for port translation pool: 50331648
Total pools: 1
Pool                 Address                 Routing              PAT  Total
Name                 Range                   Instance                  Address
src_v4_pool-namelength-is-now-changed-upto-63-characters-length 192.0.2.1-192.0.2.1
 -                   192.0.2.2-192.0.2.2
 -                   192.0.2.3-192.0.2.3
 -                   192.0.2.4-192.0.2.4
 -                   192.0.2.5-192.0.2.5
 -                   192.0.2.6-192.0.2.6
 -                   192.0.2.7-192.0.2.7
 -                   192.0.2.8-192.0.2.8
 -                   192.0.2.9-192.0.2.9
 -                   192.0.2.10-192.0.2.10
 -                   192.0.2.11-192.0.2.11
 -                   192.0.2.12-192.0.2.12
 -                   192.0.2.13-192.0.2.13
 -                   192.0.2.14-192.0.2.14
 -                   192.0.2.15-192.0.2.15
 -                   192.0.2.16-192.0.2.16
 -                   192.0.2.17-192.0.2.17
 -                   192.0.2.18-192.0.2.18
 -                   192.0.2.19-192.0.2.19
 -                   192.0.2.20-192.0.2.20
 -                   192.0.2.21-192.0.2.21
 -                   192.0.2.22-192.0.2.22
 -                   192.0.2.23-192.0.2.23
 -                   192.0.2.24-192.0.2.24
 -                   192.0.2.25-192.0.2.25
 -                   192.0.2.26-192.0.2.26
 -                   192.0.2.27-192.0.2.27
 -                   192.0.2.28-192.0.2.28
 -                   192.0.2.29-192.0.2.29
 -                   192.0.2.30-192.0.2.30
 -                   192.0.2.31-192.0.2.31
 -                   192.0.2.32-192.0.2.32
 -                   192.0.2.33-192.0.2.33
```

```
 -                  192.0.2.34-192.0.2.34
 -                  192.0.2.35-192.0.2.35
 -                  192.0.2.36-192.0.2.36
 -                  192.0.2.37-192.0.2.37
 -                  192.0.2.38-192.0.2.38
 -                  192.0.2.39-192.0.2.39
 -                  192.0.2.40-192.0.2.40
 -                  192.0.2.41-192.0.2.41
 -                  192.0.2.42-192.0.2.42
 -                  192.0.2.43-192.0.2.43
 -                  192.0.2.44-192.0.2.44
 -                  192.0.2.45-192.0.2.45
 -                  192.0.2.46-192.0.2.46
 -                  192.0.2.47-192.0.2.47
 -                  192.0.2.48-192.0.2.48
 -                  192.0.2.49-192.0.2.49
 -                  192.0.2.50-192.0.2.50
 -                  192.0.2.51-192.0.2.51
 -                  192.0.2.52-192.0.2.52
 -                  192.0.2.53-192.0.2.53
 -                  192.0.2.54-192.0.2.54
 -                  192.0.2.55-192.0.2.55
 -                  192.0.2.56-192.0.2.56
 -                  192.0.2.57-192.0.2.57
 -                  192.0.2.58-192.0.2.58
 -                  192.0.2.59-192.0.2.59
 -                  192.0.2.60-192.0.2.60
 -                  192.0.2.61-192.0.2.61
 -                  192.0.2.62-192.0.2.62
 -                  192.0.2.63-192.0.2.63
 -                  192.0.2.64-192.0.2.64  default             yes  64

Total rules: 1
Rule name : rule_namelength_is_now_changed_upto_63_characters_length
    Rule set  : ruleset_namelength_now_changed_upto_63_characters_length
    Action    : src_v4_pool-namelength-is-now-changed-upto-63-characters-length
    From      : trust                To : untrust
```

show services nat source summary

```
user@host> show services nat source summary
Total port number usage for port translation pool: 128
Maximum port number for port translation pool: 2576980378
Total pools: 1
Interface: vms-4/0/0 , Service set: ss1
Pool Name            Address Range            PAT  Total Address
NAT_POOL1--WEB-MX-SPC3NAT_WEB_MX_SPC3NAT_web-mx--spc3_NAT-POOL 203.1.1.0-203.1.1.0
 -                   203.2.1.0-203.2.1.0
 -                   203.3.1.0-203.3.1.0
 -                   203.4.1.0-203.4.1.0
 -                   203.5.1.0-203.5.1.0
 -                   203.6.1.0-203.6.1.0
 -                   203.7.1.0-203.7.1.0
 -                   203.8.1.0-203.8.1.0
 -                   203.9.1.0-203.9.1.0
 -                   203.10.1.0-203.10.1.0
 -                   203.11.1.0-203.11.1.0
 -                   203.12.1.0-203.12.1.0
 -                   203.13.1.0-203.13.1.0
 -                   203.14.1.0-203.14.1.0
 -                   203.15.1.0-203.15.1.0
 -                   203.16.1.0-203.16.1.0
 -                   203.17.1.0-203.17.1.0
 -                   203.18.1.0-203.18.1.0
 -                   203.19.1.0-203.19.1.0
 -                   203.20.1.0-203.20.1.0
 -                   203.21.1.0-203.21.1.0
 -                   203.22.1.0-203.22.1.0
 -                   203.23.1.0-203.23.1.0
 -                   203.24.1.0-203.24.1.0
 -                   203.25.1.0-203.25.1.0
 -                   203.26.1.0-203.26.1.0
 -                   203.27.1.0-203.27.1.0
 -                   203.28.1.0-203.28.1.0
 -                   203.29.1.0-203.29.1.0
 -                   203.30.1.0-203.30.1.0
 -                   203.31.1.0-203.31.1.0
 -                   203.32.1.0-203.32.1.0
 -                   203.33.1.0-203.33.1.0
 -                   203.34.1.0-203.34.1.0
```

```
  -                   203.35.1.0-203.35.1.0
  -                   203.36.1.0-203.36.1.0
  -                   203.37.1.0-203.37.1.0
  -                   203.38.1.0-203.38.1.0
  -                   203.39.1.0-203.39.1.0
  -                   203.40.1.0-203.40.1.0
  -                   203.41.1.0-203.41.1.0
  -                   203.42.1.0-203.42.1.0
  -                   203.43.1.0-203.43.1.0
  -                   203.44.1.0-203.44.1.0
  -                   203.45.1.0-203.45.1.0
  -                   203.46.1.0-203.46.1.0
  -                   203.47.1.0-203.47.1.0
  -                   203.48.1.0-203.48.1.0
  -                   203.49.1.0-203.49.1.0
  -                   203.50.1.0-203.50.1.0
  -                   203.51.1.0-203.51.1.0
  -                   203.52.1.0-203.52.1.0
  -                   203.53.1.0-203.53.1.0
  -                   203.54.1.0-203.54.1.0
  -                   203.55.1.0-203.55.1.0
  -                   203.56.1.0-203.56.1.0
  -                   203.57.1.0-203.57.1.0
  -                   203.58.1.0-203.58.1.0
  -                   203.59.1.0-203.59.1.0
  -                   203.60.1.0-203.60.1.0
  -                   203.61.1.0-203.61.1.0
  -                   203.62.1.0-203.62.1.0
  -                   203.63.1.0-203.63.1.0
  -                   203.64.1.0-203.64.1.0    yes  64

Total rules: 1
Interface: vms-4/0/0 , Service set: ss1
Rule name : NAT_RULE1--WEB-MX-SPC3NAT_WEB_MX_SPC3NAT_web-mx--spc3_NAT-RULE
    Rule set  : NAT_RULE-SET1--WEB-MX-SPC3NAT_WEB_MX_SPC3NAT_web-mx--spc3_RSET
    Action    : NAT_POOL1--WEB-MX-SPC3NAT_WEB_MX_SPC3NAT_web-mx--spc3_NAT-POOL
```

Release Information

Command introduced in Junos OS Release 9.2.

Support for IPv6 logical systems added in Junos OS Release 12.1X45-D10.

The tenant option is introduced in Junos OS Release 18.3R1.

Release History Table

Release	Description
12.3X48-D55	Starting in Junos OS Release 12.3X48-D55, and Junos OS Release 15.1X49-D90, and Junos OS Release 17.3R1, the total number of addresses that are in use for pools with IPv6 prefixes is shown as zero (0).

RELATED DOCUMENTATION

pool (Security Source NAT)

rule (Security Source NAT)

show security nat static rule

IN THIS SECTION

- Syntax | **1246**
- Description | **1246**
- Options | **1246**
- Required Privilege Level | **1246**
- Output Fields | **1247**
- Sample Output | **1248**
- Sample Output | **1249**
- Sample Output | **1250**
- Sample Output | **1250**
- Release Information | **1251**

Syntax

```
show security nat static rule
rule-name
all
logical-system (logical-system-name )
root-logical-system
tenant (tenant-name )
```

Description

Display information about the specified static Network Address Translation (NAT) rule. Traffic directions allows you to specify from interface, from zone, or from routing-instance and packet information can be source addresses and ports, and destination addresses and ports.

Options

rule-name	Name of the rule.
all	Display information about all the static NAT rules.
logical-system	Display information about the static NAT rules for a specified logical system. Specify `all` to display information for all logical systems.
root-logical-system	Display information about the static NAT rules for the primary (root) logical system.
tenant	Display information about the static NAT rules for a specified tenant system. Specify `all` to display information for all tenant systems.

Required Privilege Level

view

Output Fields

Table 78 on page 1247 lists the output fields for the **show security nat static rule** command. Output fields are listed in the approximate order in which they appear.

Table 78: show security nat static rule Output Fields

Field Name	Field Description
Static NAT rule	Name of the static NAT rule.
Total referenced IPv4/IPv6 ip-prefixes	Number of IP prefixes referenced in source, destination, and static NAT rules. This total includes the IP prefixes configured directly, as address names, and as address set names in the rule.
Rule-set	Name of the rule set. Currently, you can configure 8 rules within the same rule set.
Description	Description of the static NAT rule.
Rule-Id	Rule identification number.
Rule position	Position of the rule that indicates the order in which it applies to traffic.
From interface	Name of the interface from which the packets flow.
From routing instance	Name of the routing instance from which the packets flow.
From zone	Name of the zone from which the packets flow.
Destination addresses	Name of the destination addresses that match the rule.
Destination NPTv6 addr	Destination address that matches the rule.
Source addresses	Name of the source addresses that match the rule.

Table 78: show security nat static rule Output Fields *(Continued)*

Field Name	Field Description
Host addresses	Name of the host addresses that match the rule.
Netmask	Subnet IP address.
Destination NPTv6 Netmask	Subnet IPv6 address.
Host routing-instance	Name of the host routing instance.
Destination port	Destination port numbers that match the rule. The default value is any.
Source port	Source port numbers that match the rule.
Total static-nat rules	Number of static NAT rules.
Translation hits	Number of times a translation in the translation table is used for a static NAT rule.
Successful sessions	Number of successful session installations after the NAT rule is matched.
Failed sessions	Number of unsuccessful session installations after the NAT rule is matched.
Number of sessions	Number of sessions that reference the specified rule.

Sample Output

show security nat static rule

```
user@host> show security nat static rule namelength-is-now-changed-upto-63-
characterslength
```

```
 Static NAT rule: namelength-is-now-changed-upto-63-characterslength               Rule-set:
namelength-is-now-changed-upto-63-characterslength  Description               : The static rule
sta-r2 is for the sales team
  Rule-Id                    : 1
  Rule position              : 1
  From zone                  : zone9
  Destination addresses      : add3
  Host addresses             : add4
  Netmask                    : 24
  Host routing-instance      : N/A
  Translation hits           : 2
   Successful sessions      : 2
   Failed sessions          : 0
 Number of sessions         : 2
```

Sample Output

show security nat static rule all tenant namelength-is-now-changed-upto-63-characterslength

```
user@host> show security nat static rule all tenant namelength-is-now-changed-upto-63-
characterslength

 Total static-nat rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 2/0
Static NAT rule: namelength-is-now-changed-upto-63-characterslength                   Rule-
set: from_zone
  Rule-Id                    : 1
  Rule position              : 1
  From zone                  : untrust
  Source addresses           : 192.0.2.0       - 192.0.2.255
  Destination addresses      : 203.0.113.203
  Host addresses             : 192.168.1.203
  Netmask                    : 32
  Host routing-instance      : N/A
  Translation hits           : 0
    Successful sessions      : 0
```

```
    Failed sessions          : 0
  Number of sessions         : 0
```

Sample Output

show security nat static rule (IPv6)

```
user@host> show security nat static rule namelength-is-now-changed-upto-63-
characterslength

Static NAT rule: namelength-is-now-changed-upto-63-characterslength                  Rule-set:
namelength-is-now-changed-upto-63-characterslength
   Rule-Id                    : 1
   Rule position              : 1
   From zone                  : trust
   Destination NPTv6 addr     : 2001:db8::
   Destination NPTv6 Netmask  : 48
   Host addresses             : 2001:db8::3000
   Netmask                    : 48
   Host routing-instance      : N/A
   Translation hits           : 0
      Successful sessions      : 0
      Failed sessions          : 0
   Number of sessions         : 0
```

Sample Output

show security nat static rule all (SRX Series devices)

```
user@host> show security nat static rule all
Total static-nat rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 2/0
Static NAT rule: static_rule_length_can_be_configured_upto_63_characters_lengthh
  Rule set                   : static_rule-set_length_can_be_configured_upto_63_characters_len
  Rule Id                    : 1
  Rule position              : 1
```

```
  From zone                  : trust
  Destination addresses      : 198.51.100.0
  Host addresses             : 203.0.113.0
  Netmask                    : 32
  Host routing-instance      : N/A
  Translation hits           : 0
    Successful sessions      : 0
  Number of sessions         : 0
```

show services nat static rule (MX-SPC3)

```
user@host> show services nat static rule
Total static-nat rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 1/1
Interface: vms-4/0/0 , Service set: ss5
Static NAT rule: NAT_RULE12_WEB-MX-SPC3NAT_WEB_MX_SPC3NAT_web-mx--spc3_NAT-RULE
  Rule set                   : NAT_RULE-SET12_WEB-MX-SPC3NAT_WEB_MX_SPC3NAT_web-mx--spc3_RSET
  Rule Id                    : 1
  Rule position              : 1
  From zone                  : ss5-ZoneIn
  Destination addresses      : 198.51.100.7
  Host addresses             : 400::
  Netmask                    : 24
  Host routing-instance      : N/A
  Translation hits           : 1
    Successful sessions      : 1
    Failed sessions          : 0
  Number of sessions         : 0
```

Release Information

Command introduced in Junos OS Release 9.3.

The `Description` output field added in Junos OS Release 12.1.

Support for IPv6 logical systems and the `Successful sessions`, `Failed sessions`, `Number of sessions`, `Source addresses` and `Source ports` output fields added in Junos OS Release 12.1X45-D10.

The `Destination NPTv6 addr` and `Destination NPTv6 Netmask` output fields added in Junos OS Release 12.3X48-D25.

The tenant option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

rule (Security Static NAT)

show security policies

IN THIS SECTION

- Syntax | **1252**
- Description | **1253**
- Options | **1253**
- Required Privilege Level | **1254**
- Output Fields | **1254**
- Sample Output | **1260**
- Release Information | **1275**

Syntax

```
show security policies
<all-logical-systems-tenants>
<checksum>
<count>
<detail>
<from-zone zone-name>
<global>
<hit-count>
<information>
```

```
<logical-system logical-system-name>
<policy-name policy-name>
<root-logical-system>
<service-set>
<start>
<tenant tenant-name>
<to-zone zone-name>
<unknown-source-identity>
<zone-context>
```

Description

Displays a summary of all security policies configured on the device. If a particular policy is specified, display information specific to that policy. The existing show commands for displaying the policies configured with multiple tenant support are enhanced. A security policy controls the traffic flow from one zone to another zone. The security policies allow you to deny, permit, reject (deny and send a TCP RST or ICMP port unreachable message to the source host), encrypt and decrypt, authenticate, prioritize, schedule, filter, and monitor the traffic attempting to cross from one security zone to another.

Options

- `all-logical-systems-tenants`—Displays all multitenancy systems.
- `checksum`—Displays the policy information checksum.
- `count`—Displays the number of policies to show. Range is 1 through 65,535.
- `detail`—(Optional) Displays a detailed view of all of the policies configured on the device.
- `from-zone`—Displays the policy information matching the given source zone.
- `global`—(Optional) Displays the policy information about global policies.
- `hit-count`—Displays the policies hit count.
- `information`—Displays the policy information.
- `logical-system`—Displays the logical system name.
- `policy-name`—(Optional) Displays the policy information matching the given policy name.

- root-logical-system—Displays root logical system as default.
- service-set—Displays the name of the service set.
- start—Displays the policies from a given position. Range is 1 through 65,535.
- tenant—Displays the name of the tenant system.
- to-zone—Displays the policy information matching the given destination zone.
- unknown-source-identity—Displays the unknown-source-identity of a policy.
- zone-context—Displays the count of policies in each context (from-zone and to-zone).

Required Privilege Level

view

Output Fields

Table 79 on page 1254 lists the output fields for the show security policies command. Output fields are listed in the approximate order in which they appear.

Table 79: show security policies Output Fields

Field Name	Field Description
From zone	Name of the source zone.
To zone	Name of the destination zone.
Policy-name	Name of the applicable policy.
Description	Description of the applicable policy.

Table 79: show security policies Output Fields *(Continued)*

Field Name	Field Description
State	Status of the policy: • enabled: The policy can be used in the policy lookup process, which determines access rights for a packet and the action taken in regard to it. • disabled: The policy cannot be used in the policy lookup process, and therefore it is not available for access control.
Index	Internal number associated with the policy.
Sequence number	Number of the policy within a given context. For example, three policies that are applicable in a from-zoneA-to-zoneB context might be ordered with sequence numbers 1, 2, 3. Also, in a from-zoneC-to-zoneD context, four policies might have sequence numbers 1, 2, 3, 4.
Source addresses	For standard display mode, the names of the source addresses for a policy. Address sets are resolved to their individual names. For detail display mode, the names and corresponding IP addresses of the source addresses for a policy. Address sets are resolved to their individual address name-IP address pairs.
Destination addresses	Name of the destination address (or address set) as it was entered in the destination zone's address book. A packet's destination address must match this value for the policy to apply to it.
source-end-user-profile	Name of the device identity profile (referred to as end-user-profile in the CLI) that contains attributes, or characteristics of a device. Specification of the device identity profile in the source-end-user-profile field is part of the device identity feature. If a device matches the attributes specified in the profile and other security policy parameters, then the security policy's action is applied to traffic issuing from the device.
Source addresses (excluded)	Name of the source address excluded from the policy.

Table 79: show security policies Output Fields *(Continued)*

Field Name	Field Description
`Destination addresses (excluded)`	Name of the destination address excluded from the policy.
`Source identities`	One or more user roles specified for a policy.
`Applications`	Name of a preconfigured or custom application whose type the packet matches, as specified at configuration time. • `IP protocol`: The Internet protocol used by the application—for example, TCP, UDP, ICMP. • `ALG`: If an ALG is explicitly associated with the policy, the name of the ALG is displayed. If `application-protocol ignore` is configured, `ignore` is displayed. Otherwise, `0` is displayed. However, even if this command shows `ALG: 0`, ALGs might be triggered for packets destined to well-known ports on which ALGs are listening, unless ALGs are explicitly disabled or when `application-protocol ignore` is not configured for custom applications. • `Inactivity timeout`: Elapsed time without activity after which the application is terminated. • `Source port range`: The low-high source port range for the session application.
`Source identity feeds`	Name of a source identity (user name) added as match criteria
`Destination identity feeds`	Name of a destination identity (user name) added as match criteria
`Dynamic Applications`	Application identification-based Layer 7 dynamic applications.

Table 79: show security policies Output Fields *(Continued)*

Field Name	Field Description
`Destination Address Translation`	Status of the destination address translation traffic: • `drop translated`—Drop the packets with translated destination addresses. • `drop untranslated`—Drop the packets without translated destination addresses.

`Application Firewall`	An application firewall includes the following:
	• `Rule-set`—Name of the rule set.
	• `Rule`—Name of the rule.
	  • `Dynamic applications`—Name of the applications.
	  • `Dynamic application groups`—Name of the application groups.
	  • `Action`—The action taken with respect to a packet that matches the application firewall rule set. Actions include the following:
	  • `permit`
	  • `deny`
	• `Default rule`—The default rule applied when the identified application is not specified in any rules of the rule set.

Table 79: show security policies Output Fields *(Continued)*

Field Name	Field Description
Action or Action-type	• The action taken for a packet that matches the policy's tuples. Actions include the following: • permit • feed • firewall-authentication • tunnel ipsec-vpn *vpn-name* • pair-policy *pair-policy-name* • source-nat pool *pool-name* • pool-set *pool-set-name* • interface • destination-nat *name* • deny • reject • services-offload
Session log	Session log entry that indicates whether the at-create and at-close flags were set at configuration time to log session information.
Scheduler name	Name of a preconfigured scheduler whose schedule determines when the policy is active and can be used as a possible match for traffic.

Table 79: show security policies Output Fields *(Continued)*

Field Name	Field Description

`Policy statistics`	• `Input bytes`—The total number of bytes presented for processing by the device.
	  • `Initial direction`—The number of bytes presented for processing by the device from the initial direction.
	  • `Reply direction`—The number of bytes presented for processing by the device from the reply direction.
	• `Output bytes`—The total number of bytes actually processed by the device.
	  • `Initial direction`—The number of bytes from the initial direction actually processed by the device.
	  • `Reply direction`—The number of bytes from the reply direction actually processed by the device.
	• `Input packets`—The total number of packets presented for processing by the device.
	  • `Initial direction`—The number of packets presented for processing by the device from the initial direction.
	  • `Reply direction`—The number of packets presented for processing by the device from the reply direction.
	• `Output packets`—The total number of packets actually processed by the device.
	  • `Initial direction`—The number of packets actually processed by the device from the initial direction.
	  • `Reply direction`—The number of packets actually processed by the device from the reply direction.
	• `Session rate`—The total number of active and deleted sessions.
	• `Active sessions`—The number of sessions currently present because of access control lookups that used this policy.
	• `Session deletions`—The number of sessions deleted since system startup.
	• `Policy lookups`—The number of times the policy was accessed to check for a match.

Table 79: show security policies Output Fields *(Continued)*

Field Name	Field Description
dynapp-redir-profile	Displays unified policy redirect profile. See *profile(dynamic-application)*.
Per policy TCP Options	Configured syn and sequence checks, and the configured TCP MSS value for the initial direction, the reverse direction or, both.
Feed	Feeds details added in the security policy. The supported feeds are: • add-source-ip-to-feed • add-destination-ip-to-feed • add-source-identity-to-feed • add-destination-identity-to-feed

Sample Output

show security policies

```
user@host> show security policies

  From zone: trust, To zone: untrust
  Policy: p1, State: enabled, Index: 4, Sequence number: 1
    Source addresses:
    sa-1-ipv4: 198.51.100.11/24
    sa-2-ipv6: 2001:db8:a0b:12f0::1/32
    sa-3-ipv6: 2001:db8:a0b:12f0::22/32
    sa-4-wc:   203.0.113.1/255.255.0.255
    Destination addresses:
    da-1-ipv4: 10.2.2.2/24
    da-2-ipv6: 2001:db8:a0b:12f0::8/32
    da-3-ipv6: 2001:db8:a0b:12f0::9/32
    da-4-wc:   192.168.22.11/255.255.0.255
    Source identities: role1, role2, role4
    Applications: any
```

```
    Action: permit, application services, log, scheduled
    Application firewall : my_ruleset1
  Policy: p2, State: enabled, Index: 5, Sequence number: 2
    Source addresses:
    sa-1-ipv4:  198.51.100.11/24
    sa-2-ipv6:  2001:db8:a0b:12f0::1/32
    sa-3-ipv6:  2001:db8:a0b:12f0::22/32
    Destination addresses:
    da-1-ipv4:  10.2.2.2/24
    da-2-ipv6:  2001:db8:a0b:12f0::1/32
    da-3-ipv6:  2001:db8:a0b:12f0::9/32
    Source identities: role1, role4
    Applications: any
    Action: deny, scheduled
```

show security policies (Dynamic Applications)

```
user@host>show security policies

  Policy: p1, State: enabled, Index: 4, Scope Policy: 0, Sequence number: 1
    Source addresses: any
    Destination addresses: any
    Applications: any
    Dynamic Applications: junos:YAHOO
    Action: deny, log
  Policy: p2, State: enabled, Index: 5, Scope Policy: 0, Sequence number: 2
    Source addresses: any
    Destination addresses: any
    Applications: any
    Dynamic Applications: junos:web, junos:web:social-networking:facebook,
    junos:TFTP, junos:QQ
    Action: permit, log
  Policy: p3, State: enabled, Index: 6, Scope Policy: 0, Sequence number: 3
    Source addresses: any
    Destination addresses: any
    Applications: any
    Dynamic Applications: junos:HTTP, junos:SSL
    Action: permit, application services, log
```

The following example displays the output with unified policies configured.

```
user@host> show security policies

Default policy: deny-all
Pre ID default policy: permit-all
From zone: trust, To zone: untrust
  Policy: p2, State: enabled, Index: 4, Scope Policy: 0, Sequence number: 1
    Source addresses: any
    Destination addresses: any
    Applications: junos-defaults
    Dynamic Applications: junos:GMAIL, junos:FACEBOOK-CHAT
    dynapp-redir-profile: profile1
```

show security policies policy-name p2

```
user@host> show security policies policy-name p2
  Policy: p2, State: enabled, Index: 4, Scope Policy: 0, Sequence number: 1
    From zones: any
    To zones: any
    Source vrf group: any
    Destination vrf group: any
    Source addresses: any
    Destination addresses: any
    Applications: any
    Dynamic Applications: any
    Action: permit, application services, feed
```

show security policies policy-name detail

```
user@host> show security policies policy-name p2 detail

Policy: p2, action-type: permit, State: enabled, Index: 4, Scope Policy: 0
  Policy Type: Configured, global
  Sequence number: 1
  From zones:
    any
  To zones:
    any
```

```
  Source vrf group:
    any
  Destination vrf group:
    any
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination ports: [0-0]
  Dynamic Application:
    any: 0
  Per policy TCP Options: SYN check: No, SEQ check: No, Window scale: No
  Intrusion Detection and Prevention: disabled
  Unified Access Control: disabled
  Feed: add-source-ip-to-feed

user@host> show security policies policy-name p1 detail

  Policy: p1, action-type: permit, State: enabled, Index: 4, Scope Policy: 0
  Description: The policy p1 is for the sales team
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses:
    sa-1-ipv4:  198.51.100.11/24
    sa-2-ipv6:  2001:db8:a0b:12f0::1/32
    sa-3-ipv6:  2001:db8:a0b:12f0::9/32
    sa-4-wc:    203.0.113.1/255.255.0.255
  Destination addresses:
    da-1-ipv4: 192.0.2.0/24
    da-2-ipv6: 2001:db8:a0b:12f0::1/32
    da-3-ipv6: 2001:db8:a0b:12f0::9/32
    da-4-wc:   192.168.22.11/255.255.0.255
    Source identities:
        role1
        role2
        role4
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
```

```
      Source port range: [0-0]
      Destination port range: [0-0]
  Destination Address Translation: drop translated
  Application firewall :
     Rule-set: my_ruleset1
       Rule: rule1
         Dynamic Applications: junos:FACEBOOK-ACCESS, junos:YMSG
         Dynamic Application groups: junos:web, junos:chat
         Action: deny
       Default rule: permit
  Session log: at-create, at-close
  Scheduler name: sch20
  Per policy TCP Options: SYN check: No, SEQ check: No
  Policy statistics:
    Input  bytes       :               18144                 545 bps
      Initial direction:                9072                   272 bps
      Reply direction  :                9072                 272 bps
    Output bytes       :               18144                 545 bps
      Initial direction:                9072                 272 bps
      Reply direction  :                9072                 272 bps
    Input  packets     :                 216                   6 pps
      Initial direction:                 108                   3 bps
      Reply direction  :                 108                   3 bps
    Output packets     :                 216                   6 pps
      Initial direction:                 108                   3 bps
      Reply direction  :                 108                   3 bps
    Session rate       :                 108                   3 sps
    Active sessions    :                  93
    Session deletions  :                  15
    Policy lookups     :                 108
```

show security policies (Services-Offload)

```
user@host> show security policies

  Policy: p1, action-type: reject, State: enabled, Index: 4, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 1
  From zone: trust, To zone: trust
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
```

```
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  dynapp-redir-profile: profile1(1)
  Per policy TCP Options: SYN check: No, SEQ check: No, Window scale: No
```

show security policies (Device Identity)

```
user@host> show security policies
From zone: trust, To zone: untrust
       Policy: dev-id-marketing, State: enabled, Index: 5, Scope Policy: 0, Sequence number: 1
         Source addresses: any
         Destination addresses: any
         source-end-user-profile: marketing-profile
         Applications: any
         Action: permit
```

show security policies detail

```
user@host> show security policies detail

Default policy: deny-all
Policy: p1, action-type: permit, services-offload:enabled , State: enabled, Index: 4, Scope
Policy: 0
  Policy Type: Configured
  Description: The policy p1 is for the sales team
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses:
    any-ipv4(global):  0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global):  0.0.0.0/0
    any-ipv6(global): ::/0
    Source identities:
```

```
        role1
        role2
        role4
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
  Policy statistics:
    Input  bytes       :               18144                 545 bps
      Initial direction:                9072                   272 bps
      Reply direction  :                9072                 272 bps
    Output bytes       :               18144                 545 bps
      Initial direction:                9072                 272 bps
      Reply direction  :                9072                 272 bps
    Input  packets     :                 216                   6 pps
      Initial direction:                 108                   3 bps
      Reply direction  :                 108                   3 bps
    Output packets     :                 216                   6 pps
      Initial direction:                 108                   3 bps
      Reply direction  :                 108                   3 bps
    Session rate       :                 108                   3 sps
    Active sessions    :                  93
    Session deletions  :                  15
    Policy lookups     :                 108
Policy: p2, action-type: permit, services-offload:enabled , State: enabled, Index: 5, Scope
Policy: 0
  Policy Type: Configured
  Description: The policy p2 is for the sales team
  Sequence number: 1
  From zone: untrust, To zone: trust
  Source addresses:
    any-ipv4(global):  0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global):  0.0.0.0/0
    any-ipv6(global): ::/0
    Source identities:
        role1
        role2
        role4
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
```

```
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
```

The following example displays the output with unified policies configured.

```
user@host> show security policies detail

Default policy: deny-all
Pre ID default policy: permit-all
Policy: p2, action-type: reject, State: enabled, Index: 4, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Application: junos-defaults
    IP protocol: 6, ALG: 0, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination port range: [443-443]
    IP protocol: 6, ALG: 0, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination port range: [5432-5432]
    IP protocol: 6, ALG: 0, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination port range: [80-80]
    IP protocol: 6, ALG: 0, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination port range: [3128-3128]
    IP protocol: 6, ALG: 0, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination port range: [8000-8000]
    IP protocol: 6, ALG: 0, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination port range: [8080-8080]
    IP protocol: 17, ALG: 0, Inactivity timeout: 60
      Source port range: [0-0]
      Destination port range: [1-65535]
```

```
    IP protocol: 6, ALG: 0, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination port range: [443-443]
    IP protocol: 6, ALG: 0, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination port range: [5432-5432]
    IP protocol: 6, ALG: 0, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination port range: [80-80]
    IP protocol: 6, ALG: 0, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination port range: [3128-3128]
    IP protocol: 6, ALG: 0, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination port range: [8000-8000]
    IP protocol: 6, ALG: 0, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination port range: [8080-8080]
    IP protocol: 17, ALG: 0, Inactivity timeout: 60
      Source port range: [0-0]
      Destination port range: [1-65535]
  Dynamic Application:
    junos:FACEBOOK-CHAT: 10704
    junos:GMAIL: 51
  dynapp-redir-profile: profile1(1)
  Per policy TCP Options: SYN check: No, SEQ check: No, Window scale: No
```

show security policies detail (TCP Options)

```
user@host> show security policies policy-name p2 detail
node0:
--------------------------------------------------------------------------
Policy:p2, action-type:permit, State: enabled,Index: 4, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 1
  From zone: trust, To zone: trust
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
```

```
    any-ipv6(global): ::/0
  Application: junos-defaults
    IP protocol: tcp, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
    Destination port range: [80-80]
Per policy TCP Options: SYN check: No, SEQ check: No, Window scale: No
  Dynamic-application: junos:HTTP
```

show security policies policy-name (Negated Address)

```
user@host> show security policies policy-name p1
node0:
--------------------------------------------------------------------------
From zone: trust, To zone: untrust
  Policy: p1, State: enabled, Index: 4, Scope Policy: 0, Sequence number: 1
    Source addresses(excluded): as1
    Destination addresses(excluded): as2
    Applications: any
    Action: permit
```

show security policies policy-name detail (Negated Address)

```
user@host> show security policies policy-name p1 detail
node0:
--------------------------------------------------------------------------
Policy: p1, action-type: permit, State: enabled, Index: 4, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses(excluded):
    ad1(ad): 255.255.255.255/32
    ad2(ad):  198.51.100.1/24
    ad3(ad):  198.51.100.6  ~ 198.51.100.56
    ad4(ad):  192.0.2.8/24
    ad5(ad):  198.51.100.99 ~ 198.51.100.199
    ad6(ad):  203.0.113.9/24
    ad7(ad):  203.0.113.23/24
  Destination addresses(excluded):
    ad13(ad2): 198.51.100.76/24
    ad12(ad2): 198.51.100.88/24
```

```
    ad11(ad2): 192.0.2.23 ~ 192.0.2.66
    ad10(ad2): 192.0.2.93
    ad9(ad2):  203.0.113.76 ~ 203.0.113.106
    ad8(ad2):  203.0.113.199
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
```

show security policies global

```
user@host> show security policies global policy-name Pa
node0:
--------------------------------------------------------------------------
  Global policies:
  Policy: Pa, State: enabled, Index: 6, Scope Policy: 0, Sequence number: 1
    From zones: any
    To zones: any
    Source addresses: H0
    Destination addresses: H1
    Applications: junos-http
    Action: permit
```

show security policies detail tenant

```
user@host> show security policies detail tenant TN1

Default policy: deny-all
Pre ID default policy: permit-all
Policy: p1, action-type: permit, State: enabled, Index: 4, Scope Policy: 0
Policy Type: Configured
Sequence number: 1
From zone: trust, To zone: untrust
Source addresses: any
Destination addresses: any
Application: junos-ping
IP protocol: 1, ALG: 0, Inactivity timeout: 60
ICMP Information: type=255, code=0
```

```
Application: junos-telnet
IP protocol: tcp, ALG: 0, Inactivity timeout: 1800
Source port range: [0-0]
Destination port range: [23-23]
Application: app_udp
IP protocol: udp, ALG: 0, Inactivity timeout: 1800
Source port range: [0-0]
Destination port range: [5000-5000]
Application: junos-icmp6-all
IP protocol: 58, ALG: 0, Inactivity timeout: 60
ICMP Information: type=255, code=0
Per policy TCP Options: SYN check: No, SEQ check: No, Window scale: No
Session log: at-create, at-close
Policy statistics:
Input  bytes     :                   0                   0 bps
Initial direction:                   0                   0 bps
Reply direction  :                   0                   0 bps
Output bytes     :                   0                   0 bps
Initial direction:                   0                   0 bps
Reply direction  :                   0                   0 bps
Input  packets   :                   0                   0 pps
Initial direction:                   0                   0 bps
Reply direction  :                   0                   0 bps
Output packets   :                   0                   0 pps
Initial direction:                   0                   0 bps
Reply direction  :                   0                   0 bps
Session rate     :                   0                   0 sps
Active sessions  :                   0
Session deletions:                   0
Policy lookups   :                   0
```

show security policies (threat profile feeds)

```
user@host> show security policies policy-name p2
From zone: trust, To zone: untrust
  Policy: p2, State: enabled, Index: 5, Scope Policy: 0, Sequence number: 2
    Source vrf group: any
    Destination vrf group: any
    Source addresses: any
    Destination addresses: any
```

```
    Applications: any
Source identity feeds: user_feed_1, user_feed_2
Destination identity feeds: user_feed_3, user_feed_4
    Action: permit, application services, feed
```

show security policies detail (threat profile feeds)

```
user@host> show security policies policy-name p2 detail
Policy: p2, action-type: permit, State: enabled, Index: 5, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 2
  From zone: trust, To zone: untrust
  Source vrf group:
    any
  Destination vrf group:
    any
  Source addresses:
    any-ipv4(bob_addrbook_1): 0.0.0.0/0
    any-ipv6(bob_addrbook_1): ::/0
  Destination addresses:
    any-ipv4(bob_addrbook_1): 0.0.0.0/0
    any-ipv6(bob_addrbook_1): ::/0
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination ports: [0-0]
  Source identity feeds:
user_feed_1
user_feed_2
  Destination identity feeds:
user_feed_3
user_feed_4
  Per policy TCP Options: SYN check: No, SEQ check: No, Window scale: No
  Intrusion Detection and Prevention: disabled
  Unified Access Control: disabled
  Feed: add-source-ip-to-feed
  Feed: add-destination-ip-to-feed
  Feed: add-source-identity-to-feed
  Feed: add-destination-identity-to-feed
```

show security policies detail (services-offload enabled)

```
user@host> show security policies detail
Default policy: deny-all
Default policy log Profile ID: 0
Pre ID default policy: permit-all
Policy: SOF-enable, action-type: permit, services-offload:enabled , State: enabled, Index: 5,
Scope Policy: 0
Policy Type: Configured
Sequence number: 1
From zone: trust, To zone: untrust1
Source vrf group:
any
Destination vrf group:
any
Source addresses:
any-ipv4(global): 0.0.0.0/0
any-ipv6(global): ::/0
 Destination addresses:
 any-ipv4(global): 0.0.0.0/0
 any-ipv6(global): ::/0
 Application: any
 IP protocol: 0, ALG: 0, Inactivity timeout: 0
 Source port range: [0-0]
 Destination ports: [0-0]
 Dynamic Application:
 any: 0
 Source identity feeds:
 any
Destination identity feeds:
 any
Per policy TCP Options: SYN check: No, SEQ check: No, Window scale: No
```

show security policies policy-name SOF-enable

```
user@host> show security policies policy-name SOF-enable
From zone: trust, To zone: untrust1
  Policy: SOF-enable, State: enabled, Index: 5, Scope Policy: 0, Sequence number: 1, Log Profile
ID: 0
    Source vrf group: any
```

```
    Destination vrf group: any
    Source addresses: any
    Destination addresses: any
    Applications: any
    Dynamic Applications: any
    Source identity feeds: any
    Destination identity feeds: any
    Action: permit, services-offload
```

show security policies detail (services-offload disabled)

```
user@host> show security policies detail
Default policy: deny-all
Default policy log Profile ID: 0
Pre ID default policy: permit-all
Policy: SOF-disable, action-type: permit, services-offload:disabled , State: enabled, Index: 5,
Scope Policy: 0
  Policy Type: Configured
  Sequence number: 1
  From zone: trust, To zone: untrust1
  Source vrf group:
    any
  Destination vrf group:
    any
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination ports: [0-0]
  Dynamic Application:
    any: 0
  Source identity feeds:
    any
  Destination identity feeds:
```

```
    any
  Per policy TCP Options: SYN check: No, SEQ check: No, Window scale: No
```

show security policies policy-name SOF-disable

```
user@host> show security policies policy-name SOF-disable
 From zone: trust, To zone: untrust1
  Policy: SOF-disable, State: enabled, Index: 5, Scope Policy: 0, Sequence number: 1, Log
Profile ID: 0
    Source vrf group: any
    Destination vrf group: any
    Source addresses: any
    Destination addresses: any
    Applications: any
    Dynamic Applications: any
    Source identity feeds: any
    Destination identity feeds: any
    Action: permit, no-services-offload
```

Release Information

Command modified in Junos OS Release 9.2.

Support for IPv6 addresses is added in Junos OS Release 10.2.

Support for wildcard addresses is added in Junos OS Release 11.1.

Support for global policy and services offloading is added in Junos OS Release 11.4.

Support for source-identities and the `Description` output field is added in Junos OS Release 12.1.

Support for negated address added in Junos OS Release 12.1X45-D10.

The output fields for Policy Statistics expanded, and the output fields for the `global` and `policy-name` options are expanded to include from-zone and to-zone global match criteria in Junos OS Release 12.1X47-D10.

Support for the `initial-tcp-mss` and `reverse-tcp-mss` options is added in Junos OS Release 12.3X48-D20.

Output field and description for `source-end-user-profile` option is added in Junos OS Release 15.1x49-D70.

Output field and description for `dynamic-applications` option is added in Junos OS Release 15.1x49-D100.

Output field and description for `dynapp-redir-profile` option is added in Junos OS Release 18.2R1.

The `tenant` option is introduced in Junos OS Release 18.3R1.

The `<all-logical-systems-tenants>` option is introduced in Junos OS Release 18.4R1.

The `information` option is introduced in Junos OS Release 18.4R1.

The `checksum` option is introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

- *Security Policies Overview*
- *Understanding Security Policy Rules*
- *Understanding Security Policy Elements*
- *Unified Policies Configuration Overview*

show security screen statistics

IN THIS SECTION

- Syntax | **1277**
- Description | **1277**
- Options | **1277**
- Required Privilege Level | **1277**
- Output Fields | **1277**
- Sample Output | **1281**
- Sample Output | **1282**
- Sample Output | **1283**
- Sample Output | **1284**
- Sample Output | **1284**
- Sample Output | **1285**
- Release Information | **1291**

Syntax

```
show security screen statistics <zone zone-name / interface interface-name>
logical-system <logical-system-name | all>
root-logical-system
tenant <tenant-name >
```

Description

Display intrusion detection service (IDS) security screen statistics.

Options

- `zone` *zone-name*—Display screen statistics for this security zone.
- `interface` *interface-name* —Display screen statistics for this interface.
- *logical-system-name*—Display screen statistics for the named logical system.
- `root-logical-system`—(Optional) Display screen statistics for the primary logical system only.
- `tenant`—Display the name of the tenant system.

Required Privilege Level

view

Output Fields

Table 80 on page 1278 lists the output fields for the `show security screen statistics` command. Output fields are listed in the approximate order in which they appear.

Table 80: show security screen statistics Output Fields

Field Name	Field Description
ICMP flood	Internet Control Message Protocol (ICMP) flood counter. An ICMP flood typically occurs when ICMP echo requests use all resources in responding, such that valid network traffic can no longer be processed.
UDP flood	User Datagram Protocol (UDP) flood counter. UDP flooding occurs when an attacker sends IP packets containing UDP datagrams with the purpose of slowing down the resources, such that valid connections can no longer be handled.
TCP winnuke	Number of Transport Control Protocol (TCP) WinNuke attacks. WinNuke is a denial-of-service (DoS) attack targeting any computer on the Internet running Windows.
TCP port scan	Number of TCP port scans. The purpose of this attack is to scan the available services in the hopes that at least one port will respond, thus identifying a service to target.
ICMP address sweep	Number of ICMP address sweeps. An IP address sweep can occur with the intent of triggering responses from active hosts.
IP tear drop	Number of teardrop attacks. Teardrop attacks exploit the reassembly of fragmented IP packets.
TCP SYN flood	Number of TCP SYN attacks.
IP spoofing	Number of IP spoofs. IP spoofing occurs when an invalid source address is inserted in the packet header to make the packet appear to come from a trusted source.
ICMP ping of death	ICMP ping of death counter. Ping of death occurs when IP packets are sent that exceed the maximum legal length (65,535 bytes).
IP source route option	Number of IP source route attacks.

Table 80: show security screen statistics Output Fields *(Continued)*

Field Name	Field Description
TCP address sweep	Number of TCP address sweeps.
TCP land attack	Number of land attacks. Land attacks occur when an attacker sends spoofed SYN packets containing the IP address of the victim as both the destination and source IP address.
TCP SYN fragment	Number of TCP SYN fragments.
TCP no flag	Number of TCP headers without flags set. A normal TCP segment header has at least one control flag set.
IP unknown protocol	Number of IPs.
IP bad options	Number of invalid options.
IP record route option	Number of packets with the IP record route option enabled. This option records the IP addresses of the network devices along the path that the IP packet travels.
IP timestamp option	Number of IP timestamp option attacks. This option records the time (in Universal Time) when each network device receives the packet during its trip from the point of origin to its destination.
IP security option	Number of IP security option attacks.
IP loose source route option	Number of IP loose source route option attacks. This option specifies a partial route list for a packet to take on its journey from source to destination.
IP strict source route option	Number of IP strict source route option attacks. This option specifies the complete route list for a packet to take on its journey from source to destination.

Table 80: show security screen statistics Output Fields *(Continued)*

Field Name	Field Description
IP stream option	Number of stream option attacks. This option provides a way for the 16-bit SATNET stream identifier to be carried through networks that do not support streams.
ICMP fragment	Number of ICMP fragments. Because ICMP packets contain very short messages, there is no legitimate reason for ICMP packets to be fragmented. If an ICMP packet is so large that it must be fragmented, something is amiss.
ICMP large packet	Number of large ICMP packets.
TCP SYN FIN	Number of TCP SYN FIN packets.
TCP FIN no ACK	Number of TCP FIN flags without the acknowledge (ACK) flag.
Source session limit	Number of concurrent sessions that can be initiated from a source IP address.
TCP SYN-ACK-ACK proxy	Number of TCP flags enabled with SYN-ACK-ACK. To prevent flooding with SYN-ACK-ACK sessions, you can enable the SYN-ACK-ACK proxy protection screen option. After the number of connections from the same IP address reaches the SYN-ACK-ACK proxy threshold and SRX Series devices running Junos OS reject further connection requests from that IP address.
IP block fragment	Number of IP block fragments.
Destination session limit	Number of concurrent sessions that can be directed to a single destination IP address.
UDP address sweep	Number of UDP address sweeps.

Table 80: show security screen statistics Output Fields *(Continued)*

IPv6 extension header	Number of packets filtered for the defined IPv6 extension headers.

IPv6 extension hop by hop option	Number of packets filtered for the defined IPv6 hop-by-hop option types.
IPv6 extension destination option	Number of packets filtered for the defined IPv6 destination option types.
IPv6 extension header limit	Number of packets filtered for crossing the defined IPv6 extension header limit.
IPv6 malformed header	Number of IPv6 malformed headers defined for the intrusion detection service (IDS).
ICMPv6 malformed packet	Number of ICMPv6 malformed packets defined for the IDS options.

Sample Output

show security screen statistics zone scrzone

```
user@host> show security screen statistics zone scrzone
Screen statistics:
IDS attack type                            Statistics
  ICMP flood                               0
  UDP flood                                0
  TCP winnuke                              0
  TCP port scan                            91
  ICMP address sweep                       0
  TCP sweep                                0
  UDP sweep                                0
  IP tear drop                             0
  TCP SYN flood                            0
  IP spoofing                              0
  ICMP ping of death                       0
  IP source route option                   0
  TCP land attack                          0
  TCP SYN fragment                         0
  TCP no flag                              0
```

```
  IP unknown protocol                         0
  IP bad options                              0
  IP record route option                      0
  IP timestamp option                         0
  IP security option                          0
  IP loose source route option                0
  IP strict source route option               0
  IP stream option                            0
  ICMP fragment                               0
  ICMP large packet                           0
  TCP SYN FIN                                 0
  TCP FIN no ACK                              0
  Source session limit                        0
  TCP SYN-ACK-ACK proxy                       0
  IP block fragment                           0
  Destination session limit                   0
```

Sample Output

show security screen statistics zone untrust (IPv6)

```
user@host>show security screen statistics zone untrust
Screen statistics:

IDS attack type                             Statistics
  ICMP flood                                      0
  UDP  flood                                      0
  TCP  winnuke                                    0
  ……
  IPv6   extension header                          0
  IPv6   extension hop by hop option               0
  IPv6   extension destination option              0
  IPv6   extension header limit                    0
  IPv6   malformed header                          0
  ICMPv6 malformed packet                          0
```

Sample Output

show security screen statistics interface ge-0/0/3

```
user@host> show security screen statistics interface ge-0/0/3
Screen statistics:
IDS attack type                              Statistics
  ICMP flood                                 0
  UDP flood                                  0
  TCP winnuke                                0
  TCP port scan                              91
  ICMP address sweep                         0
  TCP sweep                                  0
  UDP sweep                                  0
  IP tear drop                               0
  TCP SYN flood                              0
  IP spoofing                                0
  ICMP ping of death                         0
  IP source route option                     0
  TCP land attack                            0
  TCP SYN fragment                           0
  TCP no flag                                0
  IP unknown protocol                        0
  IP bad options                             0
  IP record route option                     0
  IP timestamp option                        0
  IP security option                         0
  IP loose source route option               0
  IP strict source route option              0
  IP stream option                           0
  ICMP fragment                              0
  ICMP large packet                          0
  TCP SYN FIN                                0
  TCP FIN no ACK                             0
  Source session limit                       0
  TCP SYN-ACK-ACK proxy                      0
  IP block fragment                          0
  Destination session limit                  0
```

Sample Output

show security screen statistics interface ge-0/0/1 (IPv6)

```
user@host> show security screen statistics interface ge-0/0/1

Screen statistics:

IDS attack type                                 Statistics
  ICMP flood                                         0
  UDP flood                                          0
  ……
  IPv6 extension header                              0
  IPv6 extension hop by hop option                   0
  IPv6 extension destination option                  0
  IPv6 extension header limit                        0
  IPv6 malformed header                              0
  ICMPv6 malformed packet                            0
```

Sample Output

show security screen statistics interface ge-0/0/1 node primary

```
user@host> show security screen statistics interface ge-0/0/1 node primary
node0:
--------------------------------------------------------------------------
Screen statistics:
IDS attack type                              Statistics
  ICMP flood                                 1
  UDP flood                                  1
  TCP winnuke                                1
  TCP port scan                              1
  ICMP address sweep                         1
  TCP sweep                                  1
  UDP sweep                                  1
  IP tear drop                               1
  TCP SYN flood                              1
  IP spoofing                                1
```

```
  ICMP ping of death                      1
  IP source route option                  1
  TCP land attack                         1
  TCP SYN fragment                        1
  TCP no flag                             1
  IP unknown protocol                     1
  IP bad options                          1
  IP record route option                  1
  IP timestamp option                     1
  IP security option                      1
  IP loose source route option            1
  IP strict source route option           1
  IP stream option                        1
  ICMP fragment                           1
  ICMP large packet                       1
  TCP SYN FIN                             1
  TCP FIN no ACK                          1
  Source session limit                    1
  TCP SYN-ACK-ACK proxy                   1
  IP block fragment                       1
  Destination session limit               1
```

Sample Output

show security screen statistics zone trust logical-system all

```
user@host> show security screen statistics zone trust logical-system all
Logical system: root-logical-system
Screen statistics:

IDS attack type                         Statistics
  ICMP flood                            0
  UDP flood                             0
  TCP winnuke                           0
  TCP port scan                         0
  ICMP address sweep                    0
  TCP sweep                             0
  UDP sweep                             0
  IP tear drop                          0
```

```
  TCP SYN flood                          0
  IP spoofing                            0
  ICMP ping of death                     0
  IP source route option                 0
  TCP land attack                        0
  TCP SYN fragment                       0
  TCP no flag                            0
  IP unknown protocol                    0
  IP bad options                         0
  IP record route option                 0
  IP timestamp option                    0
  IP security option                     0
  IP loose source route option           0
  IP strict source route option          0
  IP stream option                       0
  ICMP fragment                          0
  ICMP large packet                      0
  TCP SYN FIN                            0
  TCP FIN no ACK                         0
  Source session limit                   0
  TCP SYN-ACK-ACK proxy                  0
  IP block fragment                      0
  Destination session limit              0

Logical system: ls1
Screen statistics:

IDS attack type                          Statistics
  ICMP flood                             0
  UDP flood                              0
  TCP winnuke                            0
  TCP port scan                          0
  ICMP address sweep                     0
  TCP sweep                              0
  UDP sweep                              0
  IP tear drop                           0
  TCP SYN flood                          0
  IP spoofing                            0
  ICMP ping of death                     0
  IP source route option                 0
  TCP land attack                        0
  TCP SYN fragment                       0
  TCP no flag                            0
```

```
  IP unknown protocol                       0
  IP bad options                            0
  IP record route option                    0
  IP timestamp option                       0
  IP security option                        0
  IP loose source route option              0
  IP strict source route option             0
  IP stream option                          0
  ICMP fragment                             0
  ICMP large packet                         0
  TCP SYN FIN                               0
  TCP FIN no ACK                            0
  Source session limit                      0
  TCP SYN-ACK-ACK proxy                     0
  IP block fragment                         0
  Destination session limit                 0

Logical system: ls2
Screen statistics:

IDS attack type                             Statistics
  ICMP flood                                0
  UDP flood                                 0
  TCP winnuke                               0
  TCP port scan                             0
  ICMP address sweep                        0
  TCP sweep                                 0
  UDP sweep                                 0
  IP tear drop                              0
  TCP SYN flood                             0
  IP spoofing                               0
  ICMP ping of death                        0
  IP source route option                    0
  TCP land attack                           0
  TCP SYN fragment                          0
  TCP no flag                               0
  IP unknown protocol                       0
  IP bad options                            0
  IP record route option                    0
  IP timestamp option                       0
  IP security option                        0
  IP loose source route option              0
  IP strict source route option             0
```

```
  IP stream option                        0
  ICMP fragment                           0
  ICMP large packet                       0
  TCP SYN FIN                             0
  TCP FIN no ACK                          0
  Source session limit                    0
  TCP SYN-ACK-ACK proxy                   0
  IP block fragment                       0
  Destination session limit               0
```

show security screen statistics zone trust tenant TN1

```
user@host> show security screen statistics zone trust tenant TN1

  Screen statistics:

  IDS attack type                         Statistics
  ICMP flood                              0
  UDP flood                               0
  TCP winnuke                             0
  TCP port scan                           0
  UDP port scan                           0
  ICMP address sweep                      0
  TCP sweep                               0
  UDP sweep                               0
  IP tear drop                            0
  TCP SYN flood                           0
  SYN flood source                        0
  SYN flood destination                   0
  IP spoofing                             0
  ICMP ping of death                      0
  IP source route option                  0
  TCP land attack                         0
  TCP SYN fragment                        0
  TCP no flag                             0
  IP unknown protocol                     0
  IP bad options                          0
  IP record route option                  0
  IP timestamp option                     0
  IP security option                      0
  IP loose source route option            0
```

```
  IP strict source route option              0
  IP stream option                           0
  ICMP fragment                              0
  ICMP large packet                          0
  TCP SYN FIN                                0
  TCP FIN no ACK                             0
  Source session limit                       0
  TCP SYN-ACK-ACK proxy                      0
  IP block fragment                          0
  Destination session limit                  0
  IPv6 extension header                      0
  IPv6 extension hop by hop option           0
  IPv6 extension destination option          0
  IPv6 extension header limit                0
  IPv6 malformed header                      0
  ICMPv6 malformed packet                    0
  IP tunnel summary                          0
```

show security screen statistics zone trust tenant all

```
user@host> show security screen statistics zone trust tenant all

  Logical system: root-logical-system
  creen statistics:

  IDS attack type                            Statistics
  ICMP flood                                 0
  UDP flood                                  0
  TCP winnuke                                0
  TCP port scan                              0
  UDP port scan                              0
  ICMP address sweep                         0
  TCP sweep                                  0
  UDP sweep                                  0
  IP tear drop                               0
  TCP SYN flood                              0
  SYN flood source                           0
  SYN flood destination                      0
  IP spoofing                                0
  ICMP ping of death                         0
  IP source route option                     0
```

```
TCP land attack                          0
TCP SYN fragment                         0
TCP no flag                              0
IP unknown protocol                      0
IP bad options                           0
IP record route option                   0
IP timestamp option                      0
IP security option                       0
IP loose source route option             0
IP strict source route option            0
IP stream option                         0
ICMP fragment                            0
ICMP large packet                        0
TCP SYN FIN                              0
TCP FIN no ACK                           0
Source session limit                     0
TCP SYN-ACK-ACK proxy                    0
IP block fragment                        0
Destination session limit                0
IPv6 extension header                    0
IPv6 extension hop by hop option         0
IPv6 extension destination option        0
IPv6 extension header limit              0
IPv6 malformed header                    0
ICMPv6 malformed packet                  0
IP tunnel summary                        0

Tenant: TN1
Screen statistics:

IDS attack type                          Statistics
ICMP flood                               0
UDP flood                                0
TCP winnuke                              0
TCP port scan                            0
UDP port scan                            0
ICMP address sweep                       0
TCP sweep                                0
UDP sweep                                0
IP tear drop                             0
TCP SYN flood                            0
SYN flood source                         0
SYN flood destination                    0
```

```
IP spoofing                              0
ICMP ping of death                       0
IP source route option                   0
TCP land attack                          0
TCP SYN fragment                         0
TCP no flag                              0
IP unknown protocol                      0
IP bad options                           0
IP record route option                   0
IP timestamp option                      0
IP security option                       0
IP loose source route option             0
IP strict source route option            0
IP stream option                         0
ICMP fragment                            0
ICMP large packet                        0
TCP SYN FIN                              0
TCP FIN no ACK                           0
Source session limit                     0
TCP SYN-ACK-ACK proxy                    0
IP block fragment                        0
Destination session limit                0
IPv6 extension header                    0
IPv6 extension hop by hop option         0
IPv6 extension destination option        0
IPv6 extension header limit              0
IPv6 malformed header                    0
ICMPv6 malformed packet                  0
IP tunnel summary                        0
```

Release Information

Command introduced in Junos OS Release 8.5.

The `node` option added in Junos OS Release 9.0.

The `logical-system all` option added in Junos OS Release 11.2R6.

Support for IPv6 extension header screens added in Junos OS Release 12.1X46-D10.

The tenant option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

- *clear security screen statistics*
- *clear security screen statistics interface*
- *clear security screen statistics zone*
- *Example: Configuring Multiple Screening Options*

show services user-identification logical-domain-identity-management

IN THIS SECTION

- Syntax | 1292
- Description | 1293
- Options | 1293
- Required Privilege Level | 1293
- Output Fields | 1293
- Sample Output | 1295
- Release Information | 1296

Syntax

```
show services user-identification logical-domain-identity-management
<counters>
<status>
```

Description

Displays the information about the logical domain identity-management.

Options

counters Displays the logical domain identity management query counters.

status Displays logical domain identity management query status.

Required Privilege Level

view

Output Fields

Table 81 on page 1293 lists the output fields for the `show services user-identification logical-domain-identity-management` command. Output fields are listed in the approximate order in which they appear.

Table 81: show services user-identification logical-domain-identity-management Output Fields

Field Name	Field Description
Node	Node (device) in the Juniper Identity Management Service server (node0 or node1).
Query server name	A Juniper Identity Management Service server stores the records for a domain name and responds to queries from clients based on these records.
For primary server:	
Address	For the status option, the IP address of the primary server.

Table 81: show services user-identification logical-domain-identity-management Output Fields *(Continued)*

Field Name	Field Description
Port	The name of the port.
Connection method	The SRX Series device that connects to the Juniper Identity Management Service server to obtain user identity information.
Connection status	The status of SRX series device connection to the Juniper Identity Management Service server.
Last received status message	The last message received.
Access token	Token string.
Token expiry time	The access token expiry time.
For secondary server:	
Address	For the status option, the IP address of the secondary server.
Batch query sent number	A number indicating how many batch queries the SRX Series device sent to the Juniper Identity Management Service server.
Batch query total response number	A number indicating how many responses the SRX Series device received from the Juniper Identity Management Service server in response to its batch queries.
Batch query error response number	A number indicating how many error responses the SRX Series device received from the Juniper Identity Management Service server in response to its batch queries.
Batch query last response time	Timestamp indicating when the last response was received.

Table 81: show services user-identification logical-domain-identity-management Output Fields *(Continued)*

Field Name	Field Description
IP query sent number	A number indicating how many IP queries the SRX Series device sent to the Juniper Identity Management Service server.
IP query total response number	A number indicating how many responses the SRX Series device received from the Juniper Identity Management Service server in response to its IP queries.
IP query error response number	A number indicating how many error responses the SRX Series device received from the Juniper Identity Management Service server in response to its IP queries.
IP query last response time	Timestamp indicating when the last response was received.

Sample Output

show services user-identification logical-domain-identity-management status

```
user@host> show services user-identification logical-domain-identity-management status
node0:
--------------------------------------------------------------------------
Query server name                :jims1
   Primary server :
     Address                     : 192.0.2.0/24
     Port                        : 443
     Connection method           : HTTPS
     Connection status           : Online
     Last received status message : OK (200)
     Access token                : isdHIbl8BXwxFftMRubGVsELRukYXtW3rtKmHiL
     Token expire time           : 2017-11-27 23:45:22
  Secondary server :
     Address                     : Not configured
```

show services user-identification logical-domain-identity-management counters

```
user@host> show services user-identification logical-domain-identity-management counters
node0:
--------------------------------------------------------------------------
Query server name                      :jims1
  Primary server :
     Address                           : 203.0.113.0/24
     Batch query sent number           : 65381
     Batch query total response number : 64930
     Batch query error response number : 38
     Batch query last response time    : 2018-08-14 15:10:52
     IP query sent number              : 10
     IP query total response number    : 10
     IP query error response number    : 0
     IP query last response time       : 2018-08-13 12:41:56
  Secondary server :
 Address                           : Not configured
```

Release Information

Command introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

clear services user-identification logical-domain-identity-management counters | 942

show system security-profile

IN THIS SECTION

- Syntax | 1297
- Description | 1297

- Options | 1297
- Required Privilege Level | 1299
- Output Fields | 1299
- Sample Output | 1300
- Release Information | 1304

Syntax

```
show system security-profile (all-resource | resource)
detail | terse
logical-system (logical-system-name )
root-logical-system
tenant (tenant-name )
```

Description

Display information about a resource allocated to the logical system in a security profile. For each resource specified, the number used by the logical system and the configured maximum and reserved values are displayed.

The `show system security-profile` command can be used by the primary administrator to display resource information for the primary logical system or user logical system. This command can also be used by the user logical system administrator to display resource information for a user logical system.

Options

Either specify `all-resource` to display information about all resources allocated for the logical system, or specify one of the following resources:

- address-book—Address books.
- appfw-rule-set—Application firewall rule set entries.

- appfw-rule—Application firewall rule entries.
- auth-entry—Firewall authentication entries.
- cpu—CPU utilization.
- flow-gate—Flow gates, also known as pinholes.
- flow-session—Flow sessions.
- icap-redirect-profile—ICAP redirect profile resource information.
- nat-cone-binding—Network Address Translation (NAT) cone bindings.
- nat-destination-pool—NAT destination pools.
- nat-destination-rule—NAT destination rules.
- nat-nopat-address—NAT without port address translations.
- nat-pat-address—NAT with port address translations.
- nat-pat-portnum—NAT source port numbers for port translation
- nat-port-ol-ipnumber—NAT port overloading IP numbers.
- nat-rule-referenced-prefix—NAT rule referenced IP-prefixes.
- nat-source-pool—NAT source pools.
- nat-source-rule—NAT source rules.
- nat-static-rule—NAT static rules.
- policy—Security policies.
- policy-with-count—Security policies with a count.
- scheduler—Schedulers.
- zone—Security zones.

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

The following options are available only to the primary administrator:

- logical-system—Display resource information for a specified user logical system. Specify `all` to display resource information for all logical systems, including the primary logical system.
- root-logical-system—Display resource information for the primary (root) logical system.

- summary—Display summary information about the resource for all logical systems.
- tenant—Display resource information for a specified tenant system. Specify `all` to display resource information for all tenant systems.

Required Privilege Level

view

Output Fields

Table 82 on page 1299 lists the output fields for the `show system security-profile` command. Output fields are listed in the approximate order in which they appear.

Table 82: show system security-profile Output Fields

Field Name	Field Description
`logical system name`	Name of the logical system.
`tenant name`	Name of the tenant system.
`security profile name`	Name of the security profile bound to the logical system.
`usage`	Number of resources that are currently being used by the logical system.
`reserved`	Number of resources that are guaranteed to be available to the logical system.
`maximum`	Number of resources that the logical system can use. The maximum does not guarantee that the amount specified for the resource in the security profile is available. The maximum is not applicable for CPU resources.
`CPU control`	`TRUE` if CPU control is enabled or `FALSE` if CPU control is not enabled.

Table 82: show system security-profile Output Fields *(Continued)*

Field Name	Field Description
`CPU control target`	Upper limit for CPU utilization on the device. The default value is 80 percent.
`CPU name`	Central point (CP) or services processing unit (SPU). CP utilization and average utilization of all SPUs is shown. The `detail` option shows CPU utilization on each SPU.
`drop rate`	Packets dropped for CPU control.

Sample Output

show system security-profile all-resource

```
user@host> show system security-profile all-resource

resource                                usage    reserved     maximum

[logical system name:   root-logical-system]
[security profile name: Default-Profile]
address-book                                0           0         512
auth-entry                                  0           0  2147483647
cpu on CP                               0.00%       1.00%      80.00%
cpu on SPU                              0.00%       1.00%      80.00%
flow-gate                                   0           0      524288
flow-session                                2           0     6291456
nat-cone-binding                            0           0       65536
nat-destination-pool                        0           0        4096
nat-destination-rule                        0           0        8192
nat-nopat-address                           0           0     1048576
nat-pat-address                             0           0        2048
nat-port-ol-ipnumber                        0           0           4
nat-rule-referenced-prefix                  0           0     1048576
nat-source-pool                             0           0        2048
nat-source-rule                             0           0        8192
```

```
nat-static-rule                              0           0       20480
policy                                       0           0       40000
policy-with-count                            0           0        1024
scheduler                                    0           0          64
zone                                         0           0         512
```

show system security-profile all-resource tenant all

```
user@host> show system security-profile all-resource tenant all

resource                                 usage    reserved     maximum

[logical system or tenant name: tn1]
[security profile name:         SP1]
address-book                                 0           0        2000
appfw-profile                                0           0        2048
appfw-rule                                   0           0      114688
appfw-rule-set                               0           0       57344
auth-entry                                   0           0       50000
cpu on CP                                0.00%       0.00%      80.00%
cpu on SPU                               0.00%       0.00%      80.00%
dslite-softwire-initiator                    0           0      100000
flow-gate                                    0           0      524288
flow-session                                 0           0   119537664
icap-redirect-profile                        0           0          64
nat-cone-binding                             0           0     2097152
nat-destination-pool                         0           0        8192
nat-destination-rule                         0           0       30720
nat-interface-port-ol                        0           0         256
nat-nopat-address                            0           0     4194304
nat-pat-address                              0           0     1048576
nat-pat-portnum                              0           0  2576980378
nat-port-ol-ipnumber                         0           0         128
nat-rule-referenced-prefix                   0           0     1048576
nat-source-pool                              0           0       12288
nat-source-rule                              0           0       30720
nat-static-rule                              0           0       30720
policy                                       0           0       80000
policy-with-count                            0           0        1024
scheduler                                    0           0          64
security-log-stream-number                   1           0           3
```

```
sla-policy                                0           0        1024
zone                                      0           0        2000
```

show system security-profile policy

```
user@host> show system security-profile policy
logical system name   security profile name       usage    reserved     maximum

ls-product-design     ls-design-profile              0          40          50
```

show system security-profile cpu

```
user@host> show system security-profile cpu
CPU control: TRUE
CPU control target: 80.00%
logical system name    profile name    CPU name    usage(%)    reserved(%)    drop rate(%)
root-logical-system    Default-Profile CP             0.00%          1.00%           0.00%
root-logical-system    Default-Profile SPU            0.00%          1.00%           0.00%
```

show system security-profile cpu logical-system all

```
user@host> show system security-profile cpu logical-system all
CPU control: TRUE
CPU control target: 80.00%
logical system name    profile name    CPU name    usage(%)    reserved(%)    drop rate(%)
root-logical-system    Default-Profile CP             0.00%          1.00%           0.00%
root-logical-system    Default-Profile SPU            0.00%          1.00%           0.00%
ls-product-design      ls-design-profile CP           0.00%          0.00%           0.00%
ls-product-design      ls-design-profile SPU          0.00%          0.00%           0.00%
ls-marketing-dept      ls-acct-mrkt-profile CP        0.00%          0.00%           0.00%
ls-marketing-dept      ls-acct-mrkt-profile SPU       0.00%          0.00%           0.00%
logical system name    security profile name    usage    reserved    maximum
root-logical-system    Default-Profile         67108864     0      134217728
lsys1                  profile1                193536     6000     134217728
```

show system security-profile cpu summary

```
user@host> show system security-profile cpu summary
CPU control: TRUE
CPU control target: 80.00%

CPU type                :     CP
global used amount      :  0.00%
global maximum quota    : 80.00%
global available amount : 80.00%
total logical systems   :      3
total security profiles :      3
heaviest usage / user   :  0.00%      / root-logical-system
lightest usage / user   :  0.00%      / root-logical-system

CPU type                :    SPU
global used amount      :  0.00%
global maximum quota    : 80.00%
global available amount : 80.00%
total logical systems   :      3
total security profiles :      3
heaviest usage / user   :  0.00%      / root-logical-system
lightest usage / user   :  0.00%      / root-logical-system
```

show system security-profile nat-pat-portnum

```
user@host> show system security-profile cpu nat-pat-portnum
CPU control: TRUE
CPU control target: 80.00%
logical system name    security profile name        usage(%)    reserved(%)    maximum
root-logical-system    Default-Profile CP             67108864          0           134217728
```

show system security-profile nat-pat-portnum summary

```
user@host> show system security-profile nat-pat-portnum summary
global used amount    :67302400
global maximum quota    :134217728
global available amount   :66915328
total logical systems    :2
```

```
total security profiles   :1
heaviest usage / user    :193536 / lsys1
```

show system security-profile icap-redirect-profile logical-system all

```
user@host> show system security-profile icap-redirect-profile logical-system all
logical-system tenant name   security profile name       usage    reserved     maximum

root-logical-system          Default-Profile                2           0          64
LSYS1                        SP1                            1          30          64
LSYS2                        SP2                            1          30          64
```

Release Information

Command introduced in Junos OS Release 11.2.

Support for application firewall added in Junos OS Release 11.3.

Option to display all resources for a logical system added in Junos OS Release 11.

Resource information for ports in source NAT pools with port translation added in Release Junos OS 11.4.

The tenant option is introduced in Junos OS Release 18.3R1.

The icap redirect profile option is introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

security-profile-resources | 876

show system security-profile secure-wire

IN THIS SECTION

- Syntax | 1305
- Description | 1306
- Options | 1306
- Required Privilege Level | 1306
- Output Fields | 1306
- Sample Output | 1307
- Sample Output | 1307
- Sample Output | 1308
- Sample Output | 1308
- Sample Output | 1309
- Sample Output | 1309
- Release Information | 1309

Syntax

```
show system security-profile secure-wire
<all-logical-systems-tenants>
<detail>
<logical-system (all | logical-system-name)>
<root-logical-system>
<summary>
<tenant (all | tenant-system-name)>
<terse>
```

Description

Display information about secure wire resource allocation.

Options

- none—Display information about all configured secure wire resource allocation.
- all-logical-systems-tenants—Display information about all configured logical system and tenant system.
- detail—Display detailed information about all configured logical system and tenant system.
- logical-system (all | *logical-system-name*)—Perform this operation on all logical systems or on a particular logical system.
- root-logical-system—Perform this operation for primary logical systems.
- summary—Display detailed summary of resource allocation.
- tenant (all | *tenant-system-name*)—Secure wire does not support tenant system.
- terse—Display summary of resource allocation.

Required Privilege Level

view

Output Fields

Table 83 on page 1307 lists the output fields for the `show system security-profile secure-wire` command in the approximate order in which they appear.

Table 83: show system security-profile secure-wire Output Fields

Field Name	Field Description
logical-system tenant name	Name of the logical system and tenant system.
security profile name	Name of the security profile name.
usage	Number of usage that are currently being used by the logical system.
reserved	Number of resources that are guaranteed to be available to the logical system.
maximum	Number of resources that the logical system can use. The maximum does not guarantee that the amount specified for the resource in the security profile is available. The maximum is not applicable for CPU resources.

Sample Output

show system security-profile secure-wire

```
user@host> show system security-profile secure-wire

logical-system tenant name   security profile name       usage    reserved     maximum

root-logical-system          Default-Profile                 0           0         255
```

Sample Output

show system security-profile secure-wire logical-system LSYS1

```
user@host> show system security-profile secure-wire logical-system LSYS1
```

```
logical-system tenant name   security profile name       usage    reserved     maximum

LSYS1                        prof1                           1           1         100
```

Sample Output

show system security-profile secure-wire all-logical-systems-tenants

```
user@host> show system security-profile secure-wire all-logical-systems-tenants

logical-system tenant name   security profile name       usage    reserved     maximum

root-logical-system          Default-Profile                 0           0         255
LSYS1                        prof1                           1           1         100
```

Sample Output

show system security-profile secure-wire detail

```
user@host> show system security-profile secure-wire detail

logical-system tenant name : root-logical-system
security profile name      : Default-Profile
used amount                : 0
reserved amount            : 0
maximum quota              : 255
```

Sample Output

show system security-profile secure-wire summary

```
user@host> show system security-profile secure-wire summary

global used amount            : 1
global maximum quota          : 255
global available amount       : 254
total logical systems/tenants : 2
total security profiles       : 1
heaviest usage / user         : 1     / LSYS1
lightest usage / user         : 0     / root-logical-system
```

Sample Output

show system security-profile secure-wire terse

```
user@host> show system security-profile secure-wire terse

logical-system tenant name   security profile name       usage    reserved     maximum

root-logical-system          Default-Profile                0           0         255
```

Release Information

Command introduced in Junos OS Release 12.3X48-D10.

`logical-system` option introduced in Junos OS Release 19.3R1.

RELATED DOCUMENTATION

Understanding Secure Wire on Security Devices

show system security-profile scheduler

IN THIS SECTION

- Syntax | 1310
- Description | 1311
- Options | 1311
- Required Privilege Level | 1311
- Output Fields | 1311
- Sample Output | 1313
- Sample Output | 1313
- Sample Output | 1313
- Sample Output | 1314
- Sample Output | 1314
- Release Information | 1314

Syntax

```
show system security-profile scheduler
detail
logical-system <logical-system-name>
root-logical-system
summary
tenant <tenant-name>
terse
```

Description

Displays the number of schedulers that the user logical system administrators and primary logical system administrators can configure for their logical systems if the security profile is bound to the logical systems. The existing show command for the security-profile scheduler is enhanced with tenant support.

Options

- `detail`—Displays the detailed output.
- `logical-system`—Displays the name of the logical system.
- `root-logical-system`—Displays root logical system as default.
- `summary`—Displays the summary output.
- `tenant`—Displays the name of the tenant system.
- `terse`—Displays the terse output.

Required Privilege Level

view

Output Fields

Table 84 on page 1311 lists the output fields for the `show system security-profile scheduler` command. Output fields are listed in the approximate order in which they appear.

Table 84: show security-profile scheduler Output Fields

Field Name	Field Description
`logical-system tenant name`	Name of the logical system or tenant system name.

Table 84: show security-profile scheduler Output Fields *(Continued)*

Field Name	Field Description
security profile name	Name of the security profile bound to the logical system or tenant system.
usage/used amount	Number of resources that are currently being used.
reserved amount	Number of resources that are guaranteed to be available to the logical system or the tenant system.
maximum quota	Number of resources that the logical system or the tenant system can use. The maximum does not guarantee that the amount specified for the resource in the security profile is available. The maximum is not applicable for CPU resources.
global maximum quota	If a logical system or the tenant system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems or tenant system.
global available amount	Number of resources available across all the logical systems or tenant system.
total logical systems/tenants	Total number of logical systems and tenant systems.
heaviest usage/user	Using the most security log streams with the detailed number.
lightest usage/user	Using the least security log streams with the detailed number.
total security profiles	Total number of the resources configured for the security profile.
global used amount	Number of resources used across all the logical systems or tenant systems.

Sample Output

show system security-profile scheduler

```
root@host> show system security-profile scheduler
logical-system tenant name   security profile name       usage    reserved     maximum

root-logical-system          Default-Profile                0           0         256
```

Sample Output

show system security-profile scheduler detail

```
root@host> show system security-profile scheduler detail
logical-system tenant name : root-logical-system
security profile name      : Default-Profile
used amount                : 0
reserved amount            : 0
maximum quota              : 256
```

Sample Output

show system security-profile scheduler summary

```
root@host> show system security-profile scheduler summary
global used amount            : 0
global maximum quota          : 256
global available amount       : 256
total logical systems/tenants : 2
total security profiles       : 2
heaviest usage / user         : 0     / root-logical-system ...(2 logical system & tenants)
lightest usage / user         : 0     / root-logical-system ...(2 logical system & tenants)
```

Sample Output

show system security-profile scheduler terse

```
root@host> show system security-profile scheduler terse
logical-system tenant name   security profile name       usage    reserved     maximum

root-logical-system          Default-Profile                0           0         256
```

Sample Output

show system security-profile scheduler tenant all

```
root@host> show system security-profile scheduler tenant all
logical-system tenant name   security profile name       usage    reserved     maximum

root-logical-system          Default-Profile                0           0         256
```

Release Information

Statement introduced in Junos OS Release 11.2.

The `tenant` option is added in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

show security policies | 1252

show system security-profile security-log-stream-number detail

IN THIS SECTION

- Syntax | 1315
- Description | 1315
- Options | 1315
- Required Privilege Level | 1316
- Output Fields | 1316
- Sample Output | 1317
- Release Information | 1318

Syntax

```
show system security-profile security-log-stream-number detail
logical-system (all | logical-system-name)
tenant (all | tenant-system-name)
```

Description

Display information about a resource allocated to the logical system or tenant system in a security profile with security log stream number.

Options

logical-system-name (all | *logical-system-name*) Display resource information for all logical systems, including the primary logical system or a particular logical system.

tenant (all | *tenant-system-name*)	Display resource information for all tenant systems, including the primary logical system or a particular tenant system.

Required Privilege Level

view

Output Fields

Table 85 on page 1316 lists the output fields for the `show system security-profile security-log-stream-number summary` command. Output fields are listed in the approximate order in which they appear.

Table 85: show system security-profile security-log-stream-number summary Output Fields

Field Name	Field Description
`logical system name`	Displays the logical system name
`security profile name`	Name of the security profile
`used amount`	Number of resources that are currently being used by the logical system.
reserved amount	Reserved quota that guarantees that the resource amount specified is always available to the logical system.
maximum quota	Maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

Sample Output

show system security-profile security-log-stream-number detail logical-system all

```
user@host> show system security-profile security-log-stream-number detail logical-system all
logical system name     : root-logical-system
security profile name   : Default-Profile
used amount             : 0
reserved amount         : 0
maximum quota           : 8

logical system name     : lsys0
security profile name   : lsys_profile
used amount             : 0
reserved amount         : 0
maximum quota           : 8

logical system name     : lsys1
security profile name   : lsys_profile
used amount             : 0
reserved amount         : 0
maximum quota           : 8

logical system name     : lsys2
security profile name   : lsys_profile
used amount             : 0
reserved amount         : 0
maximum quota           : 8
```

show system security-profile security-log-stream-number detail tenant all

```
user@host> show system security-profile security-log-stream-number detail tenant all
logical-system tenant name : root-logical-system
security profile name      : Default-Profile
used amount                : 1
reserved amount            : 0
maximum quota              : 8

logical-system tenant name : LSYS1
security profile name      : p1
```

```
used amount                : 1
reserved amount            : 1
maximum quota              : 2

logical-system tenant name : TN1
security profile name      : sp2
used amount                : 0
reserved amount            : 0
maximum quota              : 8
```

Release Information

Command introduced in Junos OS Release 18.2R1.

The `tenant` option introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

security-profile-resources | 876

show system security-profile security-log-stream-number

IN THIS SECTION

- Syntax | 1319
- Description | 1319
- Options | 1319
- Required Privilege Level | 1319
- Output Fields | 1320
- Sample Output | 1321

- Release Information | 1322

Syntax

```
show system security-profile security-log-stream-number
 logical-system (all |logical-system-name)
 tenant (all |tenant-system-name)
```

Description

Display information about a resource allocated to the logical system or tenant system in a security profile. This command can be used by the primary administrator to display resource information for the primary logical system or user logical system.

Options

logical-system (all |*logical-system-name*)	Display resource information for all logical systems, including the primary logical system or a particular logical system.
tenant (all |*tenant-system-name*)	Display resource information for all tenant systems, including the primary logical system or a particular tenant system.

Required Privilege Level

view

Output Fields

Table 86 on page 1320 lists the output fields for the `show system security-profile security-log-stream-number logical-system all` command. Output fields are listed in the approximate order in which they appear.

Table 86: show system security-profile security-log-stream-number logical-system all Output Fields

Field Name	Field Description
`logical system name`	Name of the logical system.
`security profile name`	Name of the security profile bound to the logical system.
`usage`	Number of resources that are currently being used by the logical system.
`reserved`	Number of resources that are guaranteed to be available to the logical system.
`maximum`	Number of resources that the logical system can use. The maximum does not guarantee that the amount specified for the resource in the security profile is available. The maximum is not applicable for CPU resources.
`root-logical-system`	Display resource information for the primary (root) logical system.
`Default-Profile`	Specify the authentication profile to use if no profile is specified.

Table 87 on page 1320 lists the output fields for the `show system security-profile security-log-stream-number tenant all` command. Output fields are listed in the approximate order in which they appear.

Table 87: show system security-profile security-log-stream-number tenant all Output Fields

Field Name	Field Description
`logical-system tenant name`	Name of the tenant system.
`security profile name`	Name of the security profile bound to the tenant system.

Table 87: show system security-profile security-log-stream-number tenant all Output Fields *(Continued)*

Field Name	Field Description
usage	Number of resources that are currently being used by the tenant system.
reserved	Number of resources that are guaranteed to be available to the tenant system.
maximum	Number of resources that the tenant system can use. The maximum resource does not guarantee that the amount specified for the resource in the security profile is available. The maximum resource is not applicable for CPU resources.
root-logical-system	Display resource information for the primary (root) logical system.
Default-Profile	Specify the authentication profile to use if no profile is specified.

Sample Output

show system security-profile security-log-stream-number logical-system all

```
user@host> show system security-profile security-log-stream-number logical-system all
logical system name   security profile name       usage    reserved     maximum
root-logical-system   Default-Profile                1           0           3
LSYS1                 sp1                            0           1           3
LSYS2                 sp2                            1           0           3
```

show system security-profile security-log-stream-number tenant all

```
user@host> show system security-profile security-log-stream-number tenant all
logical-system tenant name   security profile name       usage    reserved     maximum

root-logical-system          Default-Profile                1           0           8
```

LSYS1	p1	1	1	2
TN1	sp2	0	0	8

Release Information

Command introduced in Junos OS Release 18.2R1.

`tenant` option introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

security-profile-resources | 876

show system security-profile security-log-stream-number summary

IN THIS SECTION

- Syntax | 1323
- Description | 1323
- Options | 1323
- Required Privilege Level | 1323
- Output Fields | 1323
- Sample Output | 1324
- Release Information | 1325

Syntax

```
show system security-profile security-log-stream-number summary  (detail |terse)
```

Description

Display summary information about the resource for all logical systems.

Options

detail Display detailed output.

terse Display terse output (default).

Required Privilege Level

view

Output Fields

Table 88 on page 1323 lists the output fields for the `show system security-profile security-log-stream-number summary` command. Output fields are listed in the approximate order in which they appear.

Table 88: show system security-profile security-log-stream-number summary Output Fields

Field Name	Field Description
`global used amount`	Number of resources that are currently being used by the logical system.

Table 88: show system security-profile security-log-stream-number summary Output Fields *(Continued)*

Field Name	Field Description
`global maximum quota`	Number of resources that the logical system can use. The maximum does not guarantee that the amount specified for the resource in the security profile is available. The maximum is not applicable for CPU resources.
`global available amount`	Number of resources that are guaranteed to be available to the logical system.
`total logical systems`	Total number of logical systems
`total security profiles`	Total number of resources configured for the security profile
`heaviest usage / user`	Using the most security log streams with the detailed number
`lightest usage / user`	Using the least security log streams with the detailed number

Sample Output

show system security-profile security-log-stream-number summary

```
user@host> show system security-profile security-log-stream-number summary
global used amount      : 0
global maximum quota    : 32
global available amount : 32
total logical systems   : 1
total security profiles : 0
heaviest usage / user   : 0     / root-logical-system
lightest usage / user   : 0     / root-logical-system
```

show system security-profile security-log-stream-number summary

```
user@host> show system security-profile security-log-stream-number summary
global used amount            : 1
global maximum quota          : 100
global available amount       : 99
total logical systems/tenants : 1
total security profiles       : 2
heaviest usage / user         : 1     / root-logical-system
lightest usage / user         : 1     / TN1
```

Release Information

Command introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

security-profile-resources | 876

show security softwires

IN THIS SECTION

- Syntax | 1326
- Description | 1326
- Options | 1326
- Required Privilege Level | 1326
- Sample Output | 1326
- Release Information | 1327

Syntax

```
show security softwires <softwire-name softwire-name>
<logical-system (all | logical-system-name)>
```

Description

Display a summary of information of all the softwire concentrators and details on concentrators with specified name.

Options

softwire-name *softwire-name*	Display the details of the specified softwire concentrator.
logical-system (all | *logical-system-name*)	Display softwire information for all logical systems or for a specified logical system. This option is only available to the primary administrator.

Required Privilege Level

view

Sample Output

command-name

```
user@host> show security softwires
Softwire Name            SC Address         Status   Number of SI connected
  SC-CSSI-1                 3001::1          Connected      2
  SC-CSSI-str00             3100::1          Active         0
```

```
  SC-CSSI-str01             3101::1          Inactive       0
  SC-CSSI-str02             3001::1          Connected      2520
```

command-name

```
user@host> show security softwires softwire-name SC-CSSI-1
Name of softwire: SC-CSSI-1
   SC status: Connected
   SC address: 3001::1
   Zone: trust
   VR ID: 0
    SI Address               SI Status                  SPU
    3001::2                  Active                     spu-1
    3001::2                  Active                     spu-21
    SI number: 2
```

command-name

```
user@host> show security softwires logical-system ls-product-design
Softwire Name            SC Address         Status   Number of SI connected
  sc_1                      3000::1         Connected      1
```

Release Information

Command introduced in Junos OS Release 10.4. The `logical-system` option introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

show security softwires map-e confidentiality status

IN THIS SECTION

- Syntax | 1328
- Description | 1328
- Required Privilege Level | 1328
- Output Fields | 1328
- Sample Output | 1329
- Release Information | 1329

Syntax

```
show security softwires map-e confidentiality status
```

Description

Displays the MAP-E confidentiality status is enabled or disabled.

Required Privilege Level

Output Fields

Sample Output

show security softwires map-e confidentiality status

```
user@host> show security softwires map-e confidentiality status

Domain           : domain1
Confidentiality  : Enabled
```

Release Information

Command introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

confidentiality | 775

show security softwires map-e domain

IN THIS SECTION

- Syntax | 1330
- Description | 1330
- Required Privilege Level | 1330
- Output Fields | 1330
- Sample Output | 1330
- Release Information | 1331

Syntax

```
show security softwires map-e domain
```

Description

Displays the details of the MAP-E BMR and FMR rules that provides mapping between the IPv4 network and IPv6 network.

Required Privilege Level

admin

Output Fields

Sample Output

show security softwires map-e domain (with confidentiality enabled)

```
user@host> show security softwires map-e domain domain_1

Role            : CE
Domain Name     : domain_1
BR Address      : confidential
End User Ipv6 prefix: 2001:db6:12:3400::/56
BMR Mapping Rule :
    Rule Name            : bmr
FMR Mapping Rule :
    Rule Name            : fmr-1
    Rule Name            : fmr-2
```

```
    Rule Name            : fmr-3
    …
    Rule Name            : fmr-n
```

show security softwires map-e domain (without confidentiality enabled)

```
user@host> show security softwires map-e domain domain_1

Role                  : CE
Domain Name           : domain_1
BR Address            : 2001:db6:ffff::1/128
End User Ipv6 prefix  : 2001:db6:12:3400::/56
BMR Mapping Rule :
    Rule Name            : bmr
    Rule Ipv4 Prefix     : 6.0.0.1/32
    Rule Ipv6 Prefix     : 2001::3400::/56
    PSID offset          : 4
    PSID length          : 0
    EA bit length        : 0
    Port SetID           : 0
FMR Mapping Rule :
    Rule Name            : fmr
    Rule Ipv4 Prefix     : 9.0.0.1/32
    Rule Ipv6 Prefix     : 2002::3400::/56
    PSID offset          : 4
    PSID length          : 0
    EA bit length        : 0
    Port SetID           : 0
```

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

map-e

show security zones

IN THIS SECTION

- Syntax | 1332
- Description | 1332
- Options | 1333
- Required Privilege Level | 1333
- Output Fields | 1333
- Sample Output | 1335
- Release Information | 1341

Syntax

```
show security zones
<zone-name>
<all-logical-systems-tenants>
<detail>
<logical-system (logical-system-name | all)>
<root-logical-system>
<tenant (tenant-name | all)>
<terse>
<type (functional | security)>
```

Description

This command displays the information about the security zones. You can define a security zone, which allows you to divide the network into different segments and apply different security options to each segment.

Options

- *none*—Displays information about all the security zones configured.
- *zone-name*—(Optional) Displays information about the specified security zone.
- `all-logical-systems-tenants`—(Optional) Displays the information about the security zone of all logical systems and tenant systems.
- `detail`—(Optional) Displays the detail level information about the security zone.
- `logical-system` *logical-system-name*—(Optional) Displays the information about the security zones of a specified logical system.
- `logical-system all`—(Optional) Displays the information about the security zones of all logical systems.
- `root-logical-system`—(Optional) Displays the information about the security zones of the root logical system.
- `tenant` *tenant-name*—(Optional) Displays the information about the security zones of a specified tenant system.
- `tenant all`—(Optional) Displays the information about the security zones of all tenant systems.
- `terse`—(Optional) Displays the specified level information about the security zone.
- `type functional`—(Optional) Displays the information for functional zones.
- `type security`—(Optional) Displays the information for security zones.

Required Privilege Level

view

Output Fields

Table 89 on page 1334 lists the output fields for the `show security zones` command. Output fields are listed in the approximate order in which they appear.

Table 89: show security zones Output Fields

Field Name	Field Description	Level of Output
`Functional zone`	Name of the functional zone.	none
`Security zone`	Name of the security zone.	detail none
`Description`	Description of the security zone.	detail none
`Policy configurable`	Whether the policy can be configured or not.	detail none
`Interfaces bound`	Number of interfaces in the zone.	detail none
`Interfaces`	List of the interfaces in the zone.	detail none
`Zone`	Name of the zone.	terse
`Type`	Type of the zone.	terse
`Logical system`	Name of the logical system.	detail
`Tenant`	Name of the tenant system.	detail

Sample Output

show security zones

```
user@host> show security zones
Functional zone: management
   Description: This is the management zone.
   Policy configurable: No
   Interfaces bound: 1
   Interfaces:
     ge-0/0/0.0
Security zone: Host
   Description: This is the host zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     fxp0.0
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
Security zone: def
   Description: This is the def zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/2.0
```

show security zones abc

```
user@host> show security zones abc
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
```

```
    Interfaces bound: 1
    Interfaces:
      ge-0/0/1.0
```

show security zones all-logical-systems-tenants

```
user@host> show security zones all-logical-systems-tenants

Logical system: root-logical-system

Security zone: HOST
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: all
  Interfaces:

Security zone: junos-host
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:

Logical system: LSYS1

Security zone: z1
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:

Logical system: TSYS1

Security zone: z3a
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:

Security zone: z3b
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
```

```
  Interfaces bound: 0
  Interfaces:

Security zone: z3c
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:
```

show security zones abc detail

```
user@host> show security zones abc detail
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
```

show security zones logical-system LSYS1

```
user@host> show security zones logical-system LSYS1

Security zone: z1
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:
```

show security zones logical-system all

```
user@host> show security zones logical-system all

Logical system: root-logical-system

Security zone: HOST
  Send reset for non-SYN session TCP packets: Off
```

```
  Policy configurable: Yes
  Interfaces bound: all
  Interfaces:

Security zone: junos-host
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:

Logical system: LSYS1

Security zone: z1
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:
```

show security zones root-logical-system

```
user@host> show security zones root-logical-system

Security zone: HOST
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: all
  Interfaces:

Security zone: junos-host
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:
```

show security zones tenant TSYS1

```
user@host> show security zones tenant TSYS1

Security zone: z3a
  Send reset for non-SYN session TCP packets: Off
```

```
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:

Security zone: z3b
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:

Security zone: z3c
 Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:
```

show security zone tenant all

```
user@host> show security zone tenant all

  Tenant: TSYS1



  Security zone: Host
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:

  Security zone: abc
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:xe-0/0/1.0

  Security zone: def
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:xe-0/0/3.0
```

show security zones terse

```
user@host> show security zones terse
Zone                    Type
my-internal             Security
my-external             Security
dmz                     Security
```

show security zones type security

```
user@host> show security zones type security

Security zone: HOST
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: all
  Interfaces:

Security zone: junos-host
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:
```

show security zones TRUST

```
user@host> show security zones TRUST
Security zone: TRUST
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:
xe-7/0/3.0
  unidirectional-session-refreshing: Yes
```

Release Information

Command introduced in Junos OS Release 8.5.

`tenant` option introduced in Junos OS Release 18.3R1.

RELATED DOCUMENTATION

Security Zones Overview

Supported System Services for Host Inbound Traffic

security-zone