

vSRX Deployment Guide for Microsoft Hyper-V

Published
2020-12-28

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.

vSRX Deployment Guide for Microsoft Hyper-V
Copyright © 2020 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 | v](#)

1

Overview

[vSRX Overview | 2](#)

[Understand vSRX with Microsoft Hyper-V | 5](#)

[Requirements for vSRX on Microsoft Hyper-V | 7](#)

[Junos OS Features Supported on vSRX | 14](#)

2

Installing vSRX in Microsoft Hyper-V

[Prepare for vSRX Deployment in Microsoft Hyper-V | 29](#)

[Deploy vSRX in a Hyper-V Host Using the Hyper-V Manager | 30](#)

[Deploy vSRX in a Hyper-V Host Using Windows PowerShell | 43](#)

3

vSRX VM Management

[Add vSRX Interfaces | 49](#)

[Add Virtual Switches | 49](#)

[Configure the vSRX to Use a VLAN | 56](#)

[Power Down a vSRX VM with Hyper-V | 59](#)

4

Configuring and Managing vSRX

[vSRX Configuration and Management Tools | 61](#)

[Configure vSRX Using the CLI | 62](#)

[Configure vSRX Using the J-Web Interface | 64](#)

[Access the J-Web Interface and Configuring vSRX | 64](#)

[Apply the Configuration | 67](#)

[Add vSRX Feature Licenses | 68](#)

[Managing Security Policies for Virtual Machines Using Junos Space Security Director | 68](#)

Software Receive Side Scaling | 69

Overview | 69

Understanding Software Receive Side Scaling Configuration | 70

Multi-Core Scaling Support with SWRSS | 71

GTP Traffic with TEID Distribution and SWRSS | 72

Overview GTP Traffic Distribution with TEID Distribution and SWRSS | 72

Enabling GTP-U TEID Distribution with SWRSS for Asymmetric Fat Tunnels | 74

5

Configuring vSRX Chassis Clusters

Configure a vSRX Chassis Cluster in Junos OS | 78

Chassis Cluster Overview | 78

Enable Chassis Cluster Formation | 79

Chassis Cluster Quick Setup with J-Web | 80

Manually Configure a Chassis Cluster with J-Web | 81

vSRX Cluster Staging and Provisioning in Hyper-V | 88

Deploying the VMs and Additional Network Adapters in Hyper-V | 88

Creating the Control Link Connection in Hyper-V | 89

Creating the Fabric Link Connection in Hyper-V | 93

Creating the Data Interfaces Using Hyper-V | 93

Prestaging the Configuration from the Console | 94

Connecting and Installing the Staging Configuration | 96

6

Troubleshooting

Finding the Software Serial Number for vSRX | 98

About This Guide

Use this guide to install the vSRX Virtual Firewall on a Microsoft Hyper-V server. This guide also includes basic vSRX configuration and management procedures.

After completing the installation and basic configuration procedures covered in this guide, refer to the Junos OS documentation for information about further software configuration.

1

CHAPTER

Overview

[vSRX Overview | 2](#)

[Understand vSRX with Microsoft Hyper-V | 5](#)

[Requirements for vSRX on Microsoft Hyper-V | 7](#)

[Junos OS Features Supported on vSRX | 14](#)

vSRX Overview

SUMMARY

In this topic you learn about vSRX architecture and its benefits.

IN THIS SECTION


- [Benefits](#) | 5

vSRX is a virtual security appliance that provides security and networking services at the perimeter or edge in virtualized private or public *cloud* environments. vSRX runs as a virtual machine (*VM*) on a standard x86 server. vSRX is built on the Junos operating system (Junos OS) and delivers networking and security features similar to those available on the software releases for the SRX Series Services Gateways.

The vSRX provides you with a complete Next-Generation Firewall (NGFW) solution, including core firewall, VPN, NAT, advanced Layer 4 through Layer 7 security services such as Application Security, intrusion detection and prevention (IPS), and UTM features including Enhanced Web Filtering and Anti-Virus. Combined with Sky ATP, the vSRX offers a cloud-based advanced anti-malware service with dynamic analysis to protect against sophisticated malware, and provides built-in machine learning to improve verdict efficacy and decrease time to remediation.

Figure 1 on page 3 shows the high-level architecture.

Figure 1: vSRX Architecture



vSRX includes the Junos control plane (JCP) and the packet forwarding engine (PFE) components that make up the data plane. vSRX uses one virtual CPU (vCPU) for the JCP and at least one vCPU for the PFE. Starting in Junos OS Release 15.1X49-D70 and Junos OS Release 17.3R1, multi-core vSRX supports scaling vCPUs and virtual RAM (vRAM). Additional vCPUs are applied to the data plane to increase performance.

Junos OS Release 18.4R1 supports a new software architecture vSRX 3.0 that removes dual OS and nested virtualization requirement of existing vSRX architecture.


In vSRX 3.0 architecture, FreeBSD 11.x is used as the guest OS and the Routing Engine and Packet Forwarding Engine runs on FreeBSD 11.x as single virtual machine for improved performance and scalability. vSRX 3.0 uses DPDK to process the data packets in the data plane. A direct Junos upgrade from vSRX to vSRX 3.0 software is not supported.

vSRX 3.0 has the following enhancements compared to vSRX:

- Removed the restriction of requiring nested VM support in hypervisors.
- Removed the restriction of requiring ports connected to control plane to have Promiscuous mode enabled.
- Improved boot time and enhanced responsiveness of the control plane during management operations.
- Improved live migration.

Figure 2 on page 4 shows the high-level architecture for vSRX 3.0

Figure 2: vSRX 3.0 Architecture



g300161

Benefits

vSRX on standard x86 servers enables you to quickly introduce new services, deliver customized services to customers, and scale security services based on dynamic needs. vSRX is ideal for public, private, and hybrid cloud environments.

Some of the key benefits of vSRX in a virtualized private or public cloud multitenant environment include:

- *Stateful firewall* protection at the tenant edge
- Faster deployment of virtual firewalls into new sites
- Ability to run on top of various hypervisors and public cloud infrastructures
- Full routing, *VPN*, core security, and networking capabilities
- Application security features (including IPS and App-Secure)
- Content security features (including Anti Virus, Web Filtering, Anti Spam, and Content Filtering)
- Centralized management with Junos Space Security Director and local management with J-Web Interface
- Juniper Networks Sky Advanced Threat Prevention (Sky ATP) integration

Release History Table

Release	Description
15.1X49-D70	Starting in Junos OS Release 15.1X49-D70 and Junos OS Release 17.3R1, multi-core vSRX supports scaling vCPUs and virtual RAM (vRAM). Additional vCPUs are applied to the data plane to increase performance.

Understand vSRX with Microsoft Hyper-V

IN THIS SECTION

- [vSRX in Microsoft Hyper-V | 6](#)

This section presents an overview of vSRX as deployed in Microsoft Hyper-V.

vSRX in Microsoft Hyper-V

Microsoft Hyper-V is a hypervisor-based virtualization technology. It provides software infrastructure and basic management tools that you can use to create and manage a virtualized server computing environment. This virtualized environment can be used to address a variety of business goals aimed at improving efficiency and reducing costs. Hyper-V works on x86- and x64-based systems running Windows.


You deploy a vSRX virtual security appliance on a Microsoft Hyper-V server to provide networking security features for the virtualized server computing environment. Hyper-V implements isolation of virtual machines in terms of a partition. The vSRX virtual machine runs in Microsoft Hyper-V as a child partition.

Note the following for deploying vSRX on a Microsoft Hyper-V server:

- Starting in Junos OS Release 15.1X49-D80 and Junos OS Release 17.3R1, you can deploy the vSRX only on Microsoft Hyper-V Server 2012 R2 or 2012.
- Starting in Junos OS Release 15.1X49-D100 and Junos OS Release 17.4R1, you can deploy the vSRX on Microsoft Hyper-V Server 2016.

Figure 3 on page 6 illustrates the deployment of a vSRX in a Hyper-V environment to provide security for applications running on one or more virtual machines.

Figure 3: vSRX Deployment in Hyper-V



Release History Table

Release	Description
15.1X49-D80	Starting in Junos OS Release 15.1X49-D80 and Junos OS Release 17.3R1, you can deploy the vSRX only on Microsoft Hyper-V Server 2012 R2 or 2012.
15.1X49-D100	Starting in Junos OS Release 15.1X49-D100 and Junos OS Release 17.4R1, you can deploy the vSRX on Microsoft Hyper-V Server 2016.

RELATED DOCUMENTATION

[Hyper-V on Windows Server 2016](#)

[Microsoft Hyper-V Overview](#)

[Microsoft Hyper-V](#)

Requirements for vSRX on Microsoft Hyper-V

IN THIS SECTION

- [Software Requirements | 7](#)
- [Hardware Requirements | 9](#)
- [Best Practices for Improving vSRX Performance | 9](#)
- [Interface Mapping for vSRX on Microsoft Hyper-V | 10](#)
- [vSRX Default Settings on Microsoft Hyper-V | 12](#)

This section presents an overview of requirements for deploying a vSRX instance on Microsoft Hyper-V.

Software Requirements

[Table 1 on page 8](#) lists the software requirements for the vSRX instance on Microsoft Hyper-V.

NOTE: Only the vSRX small flavor is supported on Microsoft Hyper-V.

Table 1: Specifications for vSRX for Microsoft Hyper-V

Component	Specification
Hypervisor support	- Starting in Junos OS Release 15.1X49-D80 and Junos OS Release 17.3R1, you can deploy the vSRX only on Microsoft Hyper-V Server 2012 R2 or 2012. - Starting in Junos OS Release 15.1X49-D100 and Junos OS Release 17.4R1, you can deploy the vSRX on Microsoft Hyper-V Server 2016.
Memory	4 GB
Disk space	16 GB (IDE or SCSI drives)
vCPUs	2
Virtual network adapters	8 Hyper-V specific network adapters

Table 2: Specifications for vSRX 3.0 for Microsoft Hyper-V

Component	Specification
Hypervisor support	Microsoft Hyper-V Windows Server 2016
Memory	4 GB
Disk space	18 GB (IDE)
vCPUs	2

Table 2: Specifications for vSRX 3.0 for Microsoft Hyper-V (Continued)

Component	Specification
Virtual network adapters	8 Hyper-V specific network adapters

Starting in Junos OS Release 19.1R1, the vSRX 3.0 instance supports guest OS with 2 vCPUs, 4-GB virtual RAM, and a 18-GB disk space on Microsoft Hyper-V and Azure for improved performance.

Hardware Requirements

[Table 3 on page 9](#) lists the hardware specifications for the host machine that runs the vSRX VM.

Table 3: Hardware Specifications for the Host Machine

Component	Specification
Host memory size	Minimum 4 GB
Host processor type	x86 or x64-based multicore processor NOTE: DPDK requires Intel Virtualization VT-x/VT-d support in the CPU. See About Intel Virtualization Technology .
Gigabit (10/100/1000baseT) Ethernet adapter	Emulates the multiport DEC 21140 10/100TX 100 MB Ethernet network adapter with one to four network connections.

Best Practices for Improving vSRX Performance

Review the following practices to improve vSRX performance.

NUMA Nodes

The x86 server architecture consists of multiple sockets and multiple cores within a socket. Each socket also has memory that is used to store packets during I/O transfers from the NIC to the host. To efficiently read packets from memory, guest applications and associated peripherals (such as the NIC) should reside within a single socket. A penalty is associated with spanning CPU sockets for memory accesses, which might result in nondeterministic performance. For vSRX, we recommend that all vCPUs for the vSRX VM are in the same physical non-uniform memory access (NUMA) node for optimal performance.



CAUTION: The Packet Forwarding Engine (PFE) on the vSRX will become unresponsive if the NUMA nodes topology is configured in the hypervisor to spread the instance's vCPUs across multiple host NUMA nodes. vSRX requires that you ensure that all vCPUs reside on the same NUMA node.

We recommend that you bind the vSRX instance with a specific NUMA node by setting NUMA node affinity. NUMA node affinity constrains the vSRX VM resource scheduling to only the specified NUMA node.

Interface Mapping for vSRX on Microsoft Hyper-V

Each network adapter defined for a vSRX is mapped to a specific interface, depending on whether the vSRX instance is a standalone VM or one of a cluster pair for high availability.

NOTE: Starting in Junos OS Release 15.1X49-D100 for vSRX, support for chassis clustering to provide network node redundancy is only available on Microsoft Hyper-V Server 2016.

Note the following:

- In standalone mode:
 - fxp0 is the out-of-band management interface.
 - ge-0/0/0 is the first traffic (revenue) interface.
- In cluster mode:
 - fxp0 is the out-of-band management interface.
 - em0 is the cluster control link for both nodes.

- Any of the traffic interfaces can be specified as the fabric links, such as ge-0/0/0 for fab0 on node 0 and ge-7/0/0 for fab1 on node 1.

Table 4 on page 11 shows the interface names and mappings for a standalone vSRX VM.

Table 4: Interface Names for a Standalone vSRX VM

Network Adapter	Interface Name in Junos OS
1	fxp0
2	ge-0/0/0
3	ge-0/0/1
4	ge-0/0/2
5	ge-0/0/3
6	ge-0/0/4
7	ge-0/0/5
8	ge-0/0/6

Table 5 on page 11 shows the interface names and mappings for a pair of vSRX VMs in a cluster (node 0 and node 1).

Table 5: Interface Names for a vSRX Cluster Pair

Network Adapter	Interface Name in Junos OS
1	fxp0 (node 0 and 1)

Table 5: Interface Names for a vSRX Cluster Pair *(Continued)*

Network Adapter	Interface Name in Junos OS
2	em0 (node 0 and 1)
3	ge-0/0/0 (node 0) ge-7/0/0 (node 1)
4	ge-0/0/1 (node 0) ge-7/0/1 (node 1)
5	ge-0/0/2 (node 0) ge-7/0/2 (node 1)
6	ge-0/0/3 (node 0) ge-7/0/3 (node 1)
7	ge-0/0/4 (node 0) ge-7/0/4 (node 1)
8	ge-0/0/5 (node 0) ge-7/0/5 (node 1)

vSRX Default Settings on Microsoft Hyper-V

vSRX requires the following basic configuration settings:

- Interfaces must be assigned IP addresses.
- Interfaces must be bound to zones.
- Policies must be configured between zones to permit or deny traffic.

[Table 6 on page 13](#) lists the factory-default settings for security policies on the vSRX.

Table 6: Factory Default Settings for Security Policies

Source Zone	Destination Zone	Policy Action
trust	untrust	permit
trust	trust	permit
untrust	trust	deny

Release History Table

Release	Description
19.1R1	Starting in Junos OS Release 19.1R1, the vSRX 3.0 instance supports guest OS with 2 vCPUs, 4-GB virtual RAM, and a 18-GB disk space on Microsoft Hyper-V and Azure for improved performance.
15.1X49-D80	Starting in Junos OS Release 15.1X49-D80 and Junos OS Release 17.3R1, you can deploy the vSRX only on Microsoft Hyper-V Server 2012 R2 or 2012.
15.1X49-D100	Starting in Junos OS Release 15.1X49-D100 and Junos OS Release 17.4R1, you can deploy the vSRX on Microsoft Hyper-V Server 2016.
15.1X49-D100	Starting in Junos OS Release 15.1X49-D100 for vSRX, support for chassis clustering to provide network node redundancy is only available on Microsoft Hyper-V Server 2016.

RELATED DOCUMENTATION

[KB Article - Interface must be in the same routing instance as the other interfaces in the zone](#)

[About Intel Virtualization Technology](#)

[DPDK Release Notes](#)

Junos OS Features Supported on vSRX

SUMMARY

This topic provides details of the Junos OS features supported and not supported on vSRX.

IN THIS SECTION

- [SRX Series Features Supported on vSRX | 14](#)
- [SRX Series Features Not Supported on vSRX | 19](#)

SRX Series Features Supported on vSRX

vSRX inherits most of the branch SRX Series features with the following considerations shown in [Table 7 on page 14](#).

To determine the Junos OS features supported on vSRX, use the Juniper Networks Feature Explorer, a Web-based application that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. Find Feature Explorer at: [Feature Explorer: vSRX](#).

Table 7: vSRX Feature Considerations

Feature	Description
IDP	The IDP feature is subscription based and must be purchased. After purchase, you can activate the IDP feature with the license key. For SRX Series IDP configuration details, see: Understanding Intrusion Detection and Prevention for SRX Series

Table 7: vSRX Feature Considerations (*Continued*)

Feature	Description	
IPSec VPNs	Starting in Junos OS Release 19.3R1, vSRX supports the following authentication algorithms and encryption algorithms: • Authentication algorithm: hmac-sha1-96 and HMAC-SHA-256-128 authentication • Encryption algorithm: aes-128-cbc Starting in Junos OS Release 20.3R1, vSRX supports 10,000 IPsec VPN tunnels. To support the increased number of IPsec VPN tunnels, a minimum of 19 vCPUs are required. Out of the 19 vCPUs, 3 vCPUs must be dedicated to RE. You must run the request system software add optional://junos-ike.tgz command the first time you wish to enable increased IPsec tunnel capacity. For subsequent software upgrades of the instance, the junos-ike package is upgraded automatically from the new Junos OS releases installed in the instance. If chassis cluster is enabled then run this command on both the nodes. You can configure the number of vCPUs allocated to Junos Routing Engine using the set security forwarding-options resource-manager cpu re <value>. NOTE: 64 G memory is required to support 10000 tunnels in PMI mode. [See show security ipsec security-associations, show security ike tunnel-map, and show security ipsec tunnel-distribution.]	
	IPsec VPN - Tunnel Scaling on vSRX	
	Types of Tunnels	Number of tunnels supported
	Site-Site VPN tunnels	2000
	AutoVPN tunnels	10,000
	IKE SA (Site-to-site)	2000
	IKE SA (AutoVPN)	10,000

Table 7: vSRX Feature Considerations (*Continued*)

Feature	Description	
	IKE SA (Site-to-site + AutoVPN)	10,000
	IPSec SA pairs (Site-to-site)	10,000 With 2000 IKE SAs, we can have 10,000 IPSec SA.
	IPSec SA pairs (AutoVPN)	10,000
	Site-to-site + AutoVPN IPSec SA pairs	2000 Site-to-site 8000 AutoVPN
	Site-to-site + AutoVPN tunnels	2000 Site-to-site 8000 AutoVPN
ISSU	ISSU is not supported.	
Logical Systems	Starting in Junos OS Release 20.1R1, you can configure logical systems and tenant systems on vSRX and vSRX 3.0 instances. With Junos OS, you can partition a single security device into multiple logical devices that can perform independent tasks. Each logical system has its own discrete administrative domain, logical interfaces, routing instances, security firewall and other security features. See Logical Systems Overview.	

Table 7: vSRX Feature Considerations (*Continued*)

Feature	Description
PowerMode IPsec	Starting in Junos OS Release 20.1R1, vSRX 3.0 instances support PowerMode IPsec that provides IPsec performance improvements using Vector Packet Processing (VPP) and Intel AES-NI instructions. PowerMode IPsec is a small software block inside the SRX PFE (SRX Packet Forwarding Engine) that is activated when PowerMode is enabled. Supported Features in PowerMode IPsec • IPsec functionality • Traffic selectors • Secure tunnel interface (st0) • All control plane IKE functionality • Auto VPN with traffic selector • Auto VPN with routing protocol • IPv6 • Stateful Layer 4 firewall • High-Availability • NAT-T Non-Supported Features in PowerMode IPsec • NAT • IPsec in IPsec • GTP/SCTP firewall • Application firewall/AppSecure • QoS • Nested tunnel • Screen

Table 7: vSRX Feature Considerations (*Continued*)

Feature	Description
	• Multicast • Host traffic
Tenant Systems	Starting in Junos OS Release 20.1R1, you can configure tenant systems on vSRX and vSRX 3.0 instances. A tenant system provides logical partitioning of the SRX device into multiple domains similar to logical systems and provides high scalability. See Tenant Systems Overview.
Transparent mode	The known behaviors for transparent mode support on vSRX are: • The default MAC learning table size is restricted to 16,383 entries. For information about configuring transparent mode for vSRX, see Layer 2 Bridging and Transparent Mode Overview.

Table 7: vSRX Feature Considerations (*Continued*)

Feature	Description
UTM	- The UTM feature is subscription based and must be purchased. After purchase, you can activate the UTM feature with the license key. - Starting in Junos OS Release 19.4R1, vSRX 3.0 instances support the Avira scan engine, which is an on-device antivirus scanning engine. See On-Device Antivirus Scan Engine. - For SRX Series UTM configuration details, see Unified Threat Management Overview. - For SRX Series UTM antispam configuration details, see Antispam Filtering Overview. Advanced resource management (vSRX 3.0)—Starting in Junos OS Release 19.4R1, vSRX 3.0 manages the additional system resource requirements for UTM-and IDP-specific services by reallocating CPU cores and extra memory. These values for memory and CPU cores are not user configured. Previously, system resources such as memory and CPU cores were fixed. You can view the allocated CPU and memory for advance security services on vSRX 3.0 instance by using the show security forward-options resource-manager settings command. To view the flow session scaling, use the show security monitoring command. [See show security monitoring and show security forward-options resource-manager settings.]

Some Junos OS software features require a license to activate the feature. To understand more about vSRX Licenses, see, [Licenses for vSRX](#). Please refer to the [Licensing Guide](#) for general information about License Management. Please refer to the product [Data Sheets](#) for further details, or contact your Juniper Account Team or Juniper Partner.

SRX Series Features Not Supported on vSRX

vSRX inherits many features from the SRX Series device product line. [Table 8 on page 20](#) lists SRX Series features that are not applicable in a virtualized environment, that are not currently supported, or that have qualified support on vSRX.

Table 8: SRX Series Features Not Supported on vSRX

SRX Series Feature	vSRX Notes
Application Layer Gateways	
Avaya H.323	Not supported
Authentication with IC Series devices	
Layer 2 enforcement in UAC deployments	Not supported NOTE: UAC-IDP and UAC-UTM also are not supported.
Chassis cluster support NOTE: Support for chassis clustering to provide network node redundancy is only available on a vSRX deployment in Contrail, VMware, KVM, and Windows Hyper-V Server 2016.	
Chassis cluster for VirtIO driver	Only supported with KVM NOTE: The link status of VirtIO interfaces is always reported as UP, so a vSRX chassis cluster cannot receive link up and link down messages from VirtIO interfaces.
Dual control links	Not supported
In-band and low-impact cluster upgrades	Not supported
LAG and LACP (Layer 2 and Layer 3)	Not supported
Layer 2 Ethernet switching	Not supported
Low-latency firewall	Not supported
Class of service	

Table 8: SRX Series Features Not Supported on vSRX *(Continued)*

SRX Series Feature	vSRX Notes
High-priority queue on SPC	Not supported
Tunnels	Only GRE and IP-IP tunnels supported NOTE: A vSRX VM deployed on Microsoft Azure Cloud does not support GRE and multicast.
Data plane security log messages (stream mode)	
TLS protocol	Not supported
Diagnostic tools	
Flow monitoring cflowd version 9	Not supported
Ping Ethernet (CFM)	Not supported
Traceroute Ethernet (CFM)	Not supported
DNS proxy	
Dynamic DNS	Not supported
Ethernet link aggregation	
LACP in standalone or chassis cluster mode	Not supported
Layer 3 LAG on routed ports	Not supported
Static LAG in standalone or chassis cluster mode	Not supported

Table 8: SRX Series Features Not Supported on vSRX *(Continued)*

SRX Series Feature	vSRX Notes
Ethernet link fault management	
Physical interface (encapsulations) • ethernet-ccc • ethernet-tcc • extended-vlan-ccc • extended-vlan-tcc	Not supported
Interface family • ccc, tcc • ethernet-switching	Not supported
Flow-based and packet-based processing	
End-to-end packet debugging	Not supported
Network processor bundling	
Services offloading	
Interfaces	
Aggregated Ethernet interface	Not supported
IEEE 802.1X dynamic VLAN assignment	Not supported
IEEE 802.1X MAC bypass	Not supported

Table 8: SRX Series Features Not Supported on vSRX *(Continued)*

SRX Series Feature	vSRX Notes
IEEE 802.1X port-based authentication control with multisuppllicant support	Not supported
Interleaving using MLFR	Not supported
PoE	Not supported
PPP interface	Not supported
PPPoE-based radio-to-router protocol	Not supported
PPPoE interface NOTE: Starting in Junos OS Release 15.1X49-D100 and Junos OS Release 17.4R1, the vSRX supports Point-to-Point Protocol over Ethernet (PPPoE) interface.	Not supported
Promiscuous mode on interfaces	Only supported if enabled on the hypervisor
IPSec and VPNs	
Acadia - Clientless VPN	Not supported
DVPN	Not supported
Hardware IPsec (bulk crypto) Cavium/RMI	Not supported
IPsec tunnel termination in routing instances	Supported on virtual router only
Multicast for AutoVPN	Not supported

Table 8: SRX Series Features Not Supported on vSRX *(Continued)*

SRX Series Feature	vSRX Notes
IPv6 support	
DS-Lite concentrator (also called Address Family Transition Router [AFTR])	Not supported
DS-Lite initiator (aka B4)	Not supported
J-Web	
Enhanced routing configuration	Not supported
New Setup wizard (for new configurations)	Not supported
PPPoE wizard	Not supported
Remote VPN wizard	Not supported
Rescue link on dashboard	Not supported
UTM configuration for Kaspersky antivirus and the default Web filtering profile	Not supported
Log file formats for system (control plane) logs	
Binary format (binary)	Not supported
WELF	Not supported
Miscellaneous	

Table 8: SRX Series Features Not Supported on vSRX *(Continued)*

SRX Series Feature	vSRX Notes
GPRS NOTE: Starting in Junos OS Release 15.1X49-D70 and Junos OS Release 17.3R1, vSRX supports GPRS.	Not supported
Hardware acceleration	Not supported
Logical systems	Not supported
Outbound SSH	Not supported
Remote instance access	Not supported
USB modem	Not supported
Wireless LAN	Not supported
MPLS	
Circuit cross-connect (CCC) and translational cross-connect (TCC)	Not supported
Layer 2 VPNs for Ethernet connections	Only if promiscuous mode is enabled on the hypervisor
Network Address Translation	
Maximize persistent NAT bindings	Not supported
Packet capture	

Table 8: SRX Series Features Not Supported on vSRX (*Continued*)

SRX Series Feature	vSRX Notes
Packet capture	Only supported on physical interfaces and tunnel interfaces, such as <i>gr</i> , <i>ip</i> , and <i>st0</i> . Packet capture is not supported on redundant Ethernet interfaces (<i>reth</i>).
Routing	
BGP extensions for IPv6	Not supported
BGP Flowspec	Not supported
BGP route reflector	Not supported
C RTP	Not supported
Switching	
Layer 3 Q-in-Q VLAN tagging	Not supported
Transparent mode	
UTM	Not supported
Unified threat management	
Express AV	Not supported
Kaspersky AV	Not supported
Upgrading and rebooting	

Table 8: SRX Series Features Not Supported on vSRX (Continued)

SRX Series Feature	vSRX Notes
Autorecovery	Not supported
Boot instance configuration	Not supported
Boot instance recovery	Not supported
Dual-root partitioning	Not supported
OS rollback	Not supported
User interfaces	
NSM	Not supported
SRC application	Not supported
Junos Space Virtual Director	Only supported with VMware

2

CHAPTER

Installing vSRX in Microsoft Hyper-V

Prepare for vSRX Deployment in Microsoft Hyper-V | 29

Deploy vSRX in a Hyper-V Host Using the Hyper-V Manager | 30

Deploy vSRX in a Hyper-V Host Using Windows PowerShell | 43


Prepare for vSRX Deployment in Microsoft Hyper-V

Note the following guidelines when deploying vSRX on a Microsoft Hyper-V server:

- Starting in Junos OS Release 15.1X49-D80 and Junos OS Release 17.3R1, you can deploy the vSRX only on Microsoft Hyper-V Server 2012 R2 or 2012.
- Starting in Junos OS Release 15.1X49-D100 and Junos OS Release 17.4R1, you can deploy the vSRX on Microsoft Hyper-V Server 2016.
- Ensure that the host CPU supports a 64-bit x86 Intel processor and is running Windows.
- Ensure that you have a user account with administrator permissions to enable the computer to deploy a vSRX virtual machine (VM) using either Microsoft Hyper-V Manager or Windows PowerShell.
- Create the virtual switches on the Hyper-V host computer necessary to support the fxp0 (out-of-band management) interface and the traffic (revenue) interface supported by the vSRX VM. You create virtual switches using either the Microsoft Hyper-V Manager or Windows PowerShell. See ["Add vSRX Interfaces" on page 49](#) for details on adding virtual switches for the vSRX VM using the Virtual Switch Manager.

[Figure 4 on page 29](#) illustrates the deployment of a vSRX in a Hyper-V environment to provide security for applications running on one or more virtual machines.

Figure 4: Example of vSRX Deployment in Hyper-V



Release History Table

Release	Description
15.1X49-D80	Starting in Junos OS Release 15.1X49-D80 and Junos OS Release 17.3R1, you can deploy the vSRX only on Microsoft Hyper-V Server 2012 R2 or 2012.
15.1X49-D100	Starting in Junos OS Release 15.1X49-D100 and Junos OS Release 17.4R1, you can deploy the vSRX on Microsoft Hyper-V Server 2016.

RELATED DOCUMENTATION

[Install Hyper-V and Create a Virtual Machine](#)
[Create a Virtual Machine in Hyper-V](#)
[Create a Virtual Switch for Hyper-V Virtual Machines](#)
[Hyper-V Virtual Switch](#)

Deploy vSRX in a Hyper-V Host Using the Hyper-V Manager

Use this procedure to deploy and configure the vSRX as a virtual security appliance in the Hyper-V environment using Hyper-V Manager.

Note the following for deploying vSRX on a Microsoft Hyper-V server:

- Starting in Junos OS Release 15.1X49-D80 and Junos OS Release 17.3R1, you can deploy the vSRX only on Microsoft Hyper-V Server 2012 R2 or 2012.
- Starting in Junos OS Release 15.1X49-D100 and Junos OS Release 17.4R1, you can deploy the vSRX on Microsoft Hyper-V Server 2016.

NOTE: To upgrade an existing vSRX instance, see *Migration, Upgrade, and Downgrade* in the *vSRX Release Notes*.

To deploy vSRX using Hyper-V Manager:

1. Download the vSRX software image for Microsoft Hyper-V from the [Juniper Networks website](#). The vSRX disk image supported by Microsoft Hyper-V is a virtual hard disk (VHD) format file.



CAUTION: Do not change the filename of the downloaded software image or the installation will fail.

2. Log onto your Hyper-V host computer using the Administrator account.
3. Open the Hyper-V Manager by selecting **Start > Administrative Tools > Hyper-V Manager**. The welcome page for Hyper-V appears the first time that you open Hyper-V Manager.
4. Create a virtual machine by selecting **Action > New > Virtual Machine**. The Before You Begin screen appears for the New Virtual Machine Wizard. Click **Next** to move through each page of the wizard, or you can click the name of a page in the left pane to move directly to that page.

5. From the Specify Name and Location page (see [Figure 5 on page 32](#)), enter a name and location for the vSRX VM that you are creating and then click **Next**. We recommend that you keep this name the same as the hostname you intend to assign to the vSRX VM.

Figure 5: Specify Name and Location Page

New Virtual Machine Wizard

Specify Name and Location

Before You Begin
Specify Name and Location
Specify Generation
Assign Memory
Configure Networking
Connect Virtual Hard Disk
Installation Options
Summary

Choose a name and location for this virtual machine.

The name is displayed in Hyper-V Manager. We recommend that you use a name that helps you easily identify this virtual machine, such as the name of the guest operating system or workload.

Name:

You can create a folder or use an existing folder to store the virtual machine. If you don't select a folder, the virtual machine is stored in the default folder configured for this server.

☐ Store the virtual machine in a different location


Location:

 If you plan to take checkpoints of this virtual machine, select a location that has enough free space. Checkpoints include virtual machine data and may require a large amount of space.

< Previous Next > Finish Cancel

6. From the Specify Generation page (see [Figure 6 on page 33](#)), keep the default setting of **Generation 1** as the generation of the vSRX VM and then click **Next**.

Figure 6: Specify Generation Page



New Virtual Machine Wizard

Specify Generation

Before You Begin
Specify Name and Location
Specify Generation
Assign Memory
Configure Networking
Connect Virtual Hard Disk
Installation Options
Summary

Choose the generation of this virtual machine.

☒ **Generation 1**
This virtual machine generation provides the same virtual hardware to the virtual machine as in previous versions of Hyper-V.

☐ **Generation 2**
This virtual machine generation provides support for features such as Secure Boot, SCSI boot, a PXE boot using a standard network adapter. Guest operating systems must be running at least Windows Server 2012 or 64-bit versions of Windows 8.

 Once a virtual machine has been created, you cannot change its generation.

< Previous Next > Finish Cancel

7. From the Assign Memory page (see [Figure 7 on page 34](#)), enter **4096 MB** as the amount of startup memory to assign to the vSRX VM. Leave **Use Dynamic Memory for this virtual machine** clear. Click **Next**.

Figure 7: Assign Memory Page

New Virtual Machine Wizard

Assign Memory

Before You Begin
Specify Name and Location
Specify Generation
Assign Memory
Configure Networking
Connect Virtual Hard Disk
Installation Options
Summary

Specify the amount of memory to allocate to this virtual machine. You can specify an amount from 3 MB through 254294 MB. To improve performance, specify more than the minimum amount recommended for the operating system.

Startup memory: MB

☐ Use Dynamic Memory for this virtual machine.


When you decide how much memory to assign to a virtual machine, consider how you intend to use the virtual machine and the operating system that it will run.

< Previous Next > Finish Cancel

8. From the Configure Networking page (see [Figure 8 on page 35](#)), select a virtual switch from a list of existing virtual switches on the Hyper-V host computer to connect to the vSRX management interface. The default is **Not connected**. Click **Next**.

NOTE: See ["Add vSRX Interfaces"](#) on page 49 for the procedure on adding virtual switches for the vSRX VM using the Virtual Switch Manager.

Figure 8: Configure Networking Page



9. From the Connect Virtual Hard Disk page (see [Figure 9 on page 36](#)), click **Use an existing virtual hard disk** and browse to the location of the vSRX virtual hard disk (VHD) file (downloaded in Step 1). Click **Next**.

Figure 9: Connect Virtual Hard Disk Page

New Virtual Machine Wizard

Connect Virtual Hard Disk

Before You Begin
Specify Name and Location
Specify Generation
Assign Memory
Configure Networking
Connect Virtual Hard Disk
Summary

A virtual machine requires storage so that you can install an operating system. You can specify the storage now or configure it later by modifying the virtual machine's properties.

☐ Create a virtual hard disk
Use this option to create a VHDX dynamically expanding virtual hard disk.

Name:
Location:
Size: GB (Maximum: 64 TB)

☒ Use an existing virtual hard disk
Use this option to attach an existing virtual hard disk, either VHD or VHDX format.


Location:

☐ Attach a virtual hard disk later
Use this option to skip this step now and attach an existing virtual hard disk later.

< Previous Next > Finish Cancel

10. After you have finished configuring the new virtual machine, verify your selections in the Summary page (see [Figure 10 on page 37](#)) and then click **Finish** to complete the installation.


Figure 10: Summary Page



11. Right-click the vSRX VM and select **Settings** from the context menu.

12. From the Settings dialog box, under the Hardware section, select **Processor**. The Processor pane appears (see [Figure 11 on page 38](#)). Enter **2** in the **Number of virtual processors** field (the default is 1).

Figure 11: Processor Pane




13. From the Settings dialog box, under the Hardware section, select **Network Adapter**. The Network Adapter pane appears (see [Figure 12 on page 40](#)).
From the Virtual switch drop-down list, select a virtual switch to assign to a network adapter to be used by the vSRX VM (see ["Add vSRX Interfaces" on page 49](#) for details on adding virtual

switches). Each network adapter that is defined for a vSRX is mapped to a specific interface. See ["Requirements for vSRX on Microsoft Hyper-V" on page 7](#) for a summary of interface names and mappings for a vSRX VM.

NOTE: If you need to add a network adapter to assign to a virtual switch, click **Add Hardware > Network Adapter > Add**.

Figure 12: Network Adapter Pane




14. Enable the MAC address spoofing function for the vSRX VM if a network adapter is to be used as an interface for Layer 2 mode support on the vSRX. From the Network Adapter pane select **Advanced Features**. The Advanced Features pane appears (see [Figure 13 on page 42](#)). Click the **Enable MAC address spoofing** check box.

MAC address spoofing allows each network adapter to change its source MAC address for outgoing packets to one that is not assigned to them. Enabling MAC address spoofing ensures those packets are not dropped by the network adapter if the source MAC address fails to match the outgoing interface MAC address.

Click **OK** when you complete your vSRX VM selections.

Figure 13: Network Adapter Advanced Features Pane



15. On Microsoft Hyper-V Server 2016, you will need to enable nested virtualization for the vSRX VM before you power on the vSRX instance. This procedure can only be performed in the Hyper-V

environment using Windows PowerShell (see, "[Deploy vSRX in a Hyper-V Host Using Windows PowerShell](#)" on page 43, Step 9). You cannot enable nested virtualization from the Hyper-V Manager because nested virtualization is not supported on Microsoft Hyper-V Server 2012.

NOTE: Nested virtualization can only be configured on a host running Microsoft Hyper-V Server 2016. In addition, Dynamic Memory must be disabled on the virtual machine containing the nested instance of Hyper-V.

16. Launch and power on the vSRX instance in the Hyper-V Manager by selecting the vSRX VM from the list of virtual machines. Right-click and select **Start** from the context menu (or select **Action > Start**).
17. Configure the basic settings for the vSRX (see "[Configure vSRX Using the CLI](#)" on page 62).

Release History Table

Release	Description
15.1X49-D80	Starting in Junos OS Release 15.1X49-D80 and Junos OS Release 17.3R1, you can deploy the vSRX only on Microsoft Hyper-V Server 2012 R2 or 2012.
15.1X49-D100	Starting in Junos OS Release 15.1X49-D100 and Junos OS Release 17.4R1, you can deploy the vSRX on Microsoft Hyper-V Server 2016.

RELATED DOCUMENTATION

[Install Hyper-V and Create a Virtual Machine](#)

[Create a Virtual Machine in Hyper-V](#)

[Virtual Machine Settings in Hyper-V Manager Explained](#)

Deploy vSRX in a Hyper-V Host Using Windows PowerShell

Use this procedure to deploy and configure the vSRX as a virtual security appliance in the Hyper-V environment using Windows PowerShell.

Note the following for deploying vSRX on a Microsoft Hyper-V server:

- Starting in Junos OS Release 15.1X49-D80 and Junos OS Release 17.3R1, you can deploy the vSRX only on Microsoft Hyper-V Server 2012 R2 or 2012.
- Starting in Junos OS Release 15.1X49-D100 and Junos OS Release 17.4R1, you can deploy the vSRX on Microsoft Hyper-V Server 2016.

NOTE: To upgrade an existing vSRX instance, see *Migration, Upgrade, and Downgrade* in the *vSRX Release Notes*.

To deploy vSRX using Windows PowerShell:

1. Download the vSRX software image for Microsoft Hyper-V from the [Juniper Networks website](#). The vSRX disk image supported by Microsoft Hyper-V is a virtual hard disk (VHD) format file.



CAUTION: Do not change the filename of the downloaded software image or the installation will fail.

2. On the Windows desktop, click the **Start** button and type **Windows PowerShell**.
3. Right-click **Windows PowerShell** and select **Run as administrator**.
4. Run the following command to enable Hyper-V using PowerShell:
Enable-WindowsOptionalFeature -Online -FeatureName Microsoft-Hyper-V -All
5. Enter the **New-VM** command to create the vSRX VM. The command syntax is as follows:
PS C:>\Users\Administrator> New-VM -Name <Name> -MemoryStartupBytes <Memory> -BootDevice <BootDevice> -VHDPATH <VHDPATH> -Path <Path> -Generation <Generation> -Switch <SwitchName>

See [Table 9 on page 44](#) for a summary of the parameters in the **New-VM** command.

Table 9: New-VM Command Parameters

Parameter	Description
<i>-Name</i>	Specify a name for the vSRX VM that you are creating. We recommend keeping this name the same as the hostname you intend to give to the vSRX VM.
<i>-MemoryStartupBytes</i>	Enter 4GB as the amount of startup memory to assign to the vSRX VM.

Table 9: New-VM Command Parameters (Continued)

Parameter	Description
<i>-BootDevice</i>	Enter VHD as the device that the vSRX VM boots to when it starts.
<i>-VHDPATH</i>	Specify the location of the vSRX virtual hard disk (VHD) file that you want to deploy.
<i>-Path</i>	Specify the location to store the vSRX VM configuration files.
<i>-Generation</i>	Enter 1 to create a generation 1 virtual machine for the vSRX.
<i>-SwitchName</i>	Specify the name of the virtual switch that you want the vSRX VM to assign to a network adapter used by the vSRX VM. Each network adapter that is defined for a vSRX is mapped to a specific interface. See "Requirements for vSRX on Microsoft Hyper-V" on page 7 for a summary of interface names and mappings for a vSRX VM. NOTE: To locate the name of a previously created virtual switch, use the Get-VMSwitch command. See "Add vSRX Interfaces" on page 49 for the procedure on adding virtual switches for the vSRX VM using the Virtual Switch Manager.

The following is an example of the **New-VM** command syntax for creating a vSRX VM:

```
PS C:>\Users\Administrator> New-VM -Name vSRX_0109 -MemoryStartupBytes 4GB -
BootDevice VHD -VHDPATH C:\Users\Public\Documents\Hyper-V\vsrx-0109-powershell\vsrx
\media-vsrx-vmdisk-151X49D80.hyper-v.vhd -Path 'C:\Users\Public\Documents\Hyper-V
\vsrx-0109\' Generation 1 SwitchName test
```

- Set the number of processors for the newly created vSRX VM by entering the **Set-VMProcessor** command. Specify **Count 2** for the number of processors. For example:

```
PS C:>\Users\Administrator> Set-VMProcessor -VMName <vSRVName> -Count 2
```

- Verify the newly created vSRX VM by entering the **Get-VM** command. For example:

```
PS C:>\Users\Administrator> Get-VM -VMName <vSRVName>
```

The output for the command is as follows:

Name	State	CPUUsage (%)	MemoryAssigned (M)	Uptime
vSRX_0109	Off	0	0	00:00:00 Operating
normally	8.0			

8. Enable the MAC address spoofing function for the vSRX VM if a network adapter is to be used as an interface for Layer 2 mode support on the vSRX. MAC address spoofing allows the vSRX VM's network adapter to change its source MAC address for outgoing packets to one that is not assigned to them. Enabling MAC address spoofing ensures those packets are not dropped by the network adapter if the source MAC address fails to match the outgoing interface MAC address.

The command syntax is as follows:

```
PS C:>\Users\Administrator> Set-VMNetworkAdapter -VMName <vSRVName> -computerName <HyperVHostName> -VMNetworkAdapter <NetworkAdapterName> -MacAddressSpoofing On
```

Verify that MacAddressSpoofing is On.

```
PS C:>\Users\Administrator> Get-VMNetworkAdapter -VMName <vSRVName> -computerName <HyperVHostName> | fl <HyperVHostName>name,macaddressspoofing
```

The output for the command is as follows:

Name	: vSRX_0109
MacAddressSpoofing	: On

9. Enable nested virtualization for the vSRX VM by using the **Set-VMProcessor** command, where **VMName** is the name of the vSRX VM you created. By default, the virtualization extensions are disabled for each VM. Nested virtualization allows you to run Hyper-V inside of a Hyper-V virtual machine. For example:

```
PS C:>\Users\Administrator> Set-VMProcessor -VMName <vSRX_0109> -ExposeVirtualizationExtensions $true
```

NOTE: Nested virtualization can only be configured on a host running Microsoft Hyper-V Server 2016. In addition, Dynamic Memory must be disabled on the virtual machine containing the nested instance of Hyper-V.

10. Launch and power on the vSRX VM by using the **Start-VM** command, where **Name** is the name of the vSRX VM you created. For example:

```
PS C:>\Users\Administrator> Start-VM -Name <vSRX_0109>
```

11. Configure the basic settings for the vSRX (see "[Configure vSRX Using the CLI](#)" on page 62).

Release History Table

Release	Description
15.1X49-D80	Starting in Junos OS Release 15.1X49-D80 and Junos OS Release 17.3R1, you can deploy the vSRX only on Microsoft Hyper-V Server 2012 R2 or 2012.
15.1X49-D100	Starting in Junos OS Release 15.1X49-D100 and Junos OS Release 17.4R1, you can deploy the vSRX on Microsoft Hyper-V Server 2016.

RELATED DOCUMENTATION

[Hyper-V Module for Windows PowerShell](#)

[Create a Virtual Machine in Hyper-V](#)

[Run Hyper-V in a Virtual Machine with Nested Virtualization](#)

3

CHAPTER

vSRX VM Management

[Add vSRX Interfaces | 49](#)

[Power Down a vSRX VM with Hyper-V | 59](#)

Add vSRX Interfaces

IN THIS SECTION

- [Add Virtual Switches | 49](#)
- [Configure the vSRX to Use a VLAN | 56](#)

The Hyper-V virtual switch is a software-based Layer 2 Ethernet network switch that connects VMs to either physical or virtual networks. A virtual switch can be configured from Hyper-V Manager or Windows PowerShell. The Hyper-V host uses the virtual switches to connect virtual machines to the internet through the host computer's network connection. You configure networking for the vSRX by adding, removing, and modifying its associated network adapters in the Hyper-V host as necessary.

NOTE: To perform this procedure, you must have appropriate permissions. Contact your Virtual Server administrator to request the proper permissions to add a virtual switch and network adapter..

For the vSRX VM, you pair a network adapter with a virtual switch for the vSRX to receive and transmit traffic. You map network adapters to the specific vSRX interfaces: Network adapter 1 is mapped to the fxp0 (out-of-band management) interface, network adapter 2 is mapped to the ge-0/0/0 (revenue) interface, network adapter 3 is mapped to ge-0/0/1, and so on (see "[Requirements for vSRX on Microsoft Hyper-V](#)" on page 7). Hyper-V supports a maximum of eight network adapters.

NOTE: When adding virtual switches, there are no limits imposed by Hyper-V. The practical limit depends on the available computing resources.

This section includes the following topics on adding vSRX interfaces in Hyper-V:


Add Virtual Switches

To add virtual switches for the vSRX VM using the Virtual Switch Manager in the Hyper-V Manager:

1. Open the Hyper-V Manager by selecting **Start > Administrative Tools > Hyper-V Manager**.

2. Select **Action > Virtual Switch Manager**. The Virtual Switch Manager appears.
3. Under the Virtual Switches section, select **New virtual network switch**. The Create Virtual Switch pane appears (see [Figure 14 on page 50](#)).

Figure 14: Create Virtual Switch Pane



4. Choose the type of virtual switch to create:

- **External**—Gives virtual machines access to a physical network to communicate with servers and clients on an external network. It allows virtual machines on the same Hyper-V server to communicate with each other.
- **Internal**—Allows communication between virtual machines on the same Hyper-V server, and between the virtual machines and the management host operating system.
- **Private**—Allows communication only between virtual machines on the same Hyper-V server. A private network is isolated from all external network traffic on the Hyper-V server. This type of network is useful when you must create an isolated networking environment, like an isolated test domain.


In most cases when adding a vSRX network adapter, select **External** as the type of virtual switch. Internal and private virtual switches are intended to keep network traffic within the Hyper-V server.

NOTE: For the fxp0 (out-of-band management) interface, connect it to External virtual switch, which could connect to an external network.

For the ge-0/0/0 (revenue port) interface, if only communication between VMs in the same Hyper-V server is needed, Internal or Private virtual switch should be sufficient. However, if communication between the VM and an external network is needed, connect it to External virtual switch.

5. Select **Create Virtual Switch**. The Virtual Switch Properties pane appears (see [Figure 15 on page 52](#)).

Figure 15: Virtual Switch Properties Pane



Virtual Switch Properties

Name:
ge-0/0/0

Notes:

Connection type
What do you want to connect this virtual switch to?

☒ External network:
Broadcom BCM5709C NetXtreme II GigE (NDIS VBD Client) #47

☒ Allow management operating system to share this network adapter

☐ Enable single-root I/O virtualization (SR-IOV)

☐ Internal network

☐ Private network

VLAN ID
☐ Enable virtual LAN identification for management operating system

The VLAN identifier specifies the virtual LAN that the management operating system will use for all network communications through this network adapter. This setting does not affect virtual machine networking.

2

Remove

i SR-IOV can only be configured when the virtual switch is created. An external virtual switch with SR-IOV enabled cannot be converted to an internal or private switch.

OK Cancel Apply

6. Specify a name for the virtual switch.
7. Choose the physical network interface card b(NIC) that you want to use (only a requirement when you select **External**).

8. Isolate network traffic from the management Hyper-V host operating system or other virtual machines that share the same virtual switch by selecting **Enable virtual LAN identification**. You can change the VLAN ID to any number or leave the default. See ["Configure the vSRX to Use a VLAN" on page 56](#) for details.
9. Click **OK**, then click **Yes** to apply networking changes and to close the Virtual Switch Manager window.
10. If necessary, repeat Steps 3 through 9 to add additional network adapters for use by the vSRX VM.
11. Right-click the vSRX VM and select **Settings** from the context menu. From the Settings dialog box, under the Hardware section, click **Network Adapter**. The Network Adapter pane appears (see [Figure 16 on page 54](#)).

12. From the Virtual switch drop-down list, select the **virtual switch** that you want to assign to this network adapter. See ["Requirements for vSRX on Microsoft Hyper-V" on page 7](#) for a summary of interface names and mappings for a vSRX VM.

Figure 16: Adding Virtual Switch to Network Adapter Example

Network Adapter

Specify the configuration of the network adapter or remove the network adapter.

Virtual switch:
ge-0/0/0

VLAN ID

☐ Enable virtual LAN identification

The VLAN identifier specifies the virtual LAN that this virtual machine will use for all network communications through this network adapter.

VLAN ID: 2

Bandwidth Management

☐ Enable bandwidth management

Specify how this network adapter utilizes network bandwidth. Both Minimum Bandwidth and Maximum Bandwidth are measured in Megabits per second.

Minimum bandwidth: 0 Mbps

Maximum bandwidth: 0 Mbps

To leave the minimum or maximum unrestricted, specify 0 as the value.

To remove the network adapter from this virtual machine, click Remove.

Remove


Use a legacy network adapter instead of this network adapter to perform a network-based installation of the guest operating system or when integration services are not installed in the guest operating system.

OK **Cancel** **Apply**

13. If a network adapter is to be used as an interface for Layer 2 mode support on the vSRX, then from the Network Adapter pane select **Advanced Features**. Select the **Enable MAC address spoofing** check box to enable the MAC address spoofing function for the network adapter (see [Figure 17 on page 55](#)).

MAC address spoofing allows each network adapter to change its source MAC address for outgoing packets to one that is not assigned to them. Enabling MAC address spoofing ensures those packets are not dropped by the network adapter if the source MAC address fails to match the outgoing interface MAC address.

Figure 17: Network Adapter Enable MAC Address Spoofing Example



14. Click **Apply** and **OK** to save the changes in the Settings dialog box.
15. Launch and power on the vSRX instance in the Hyper-V Manager by selecting the vSRX VM from the list of virtual machines, and then right-click and select **Start** from the context menu (or select **Action > Start**).

SEE ALSO

[Create a Virtual Switch for Hyper-V Virtual Machines](#)

[Create a Virtual Network](#)

Configure the vSRX to Use a VLAN

Hyper-V supports the configuration of VLANs on a network adapter in the host computer. For each network adapter that you configure for the vSRX VM, if required, you can add a VLAN identifier to specify the VLAN that the vSRX VM will use for all network communications through the network adapter.

By default, Hyper-V enables trunk mode for a VLAN. Trunk mode allows multiple VLAN IDs to share a connection between the physical network adapter and the physical network.

To give the vSRX VM external access on the virtual network in multiple VLANs, you will need to configure the port on the physical network to be in trunk mode. You will also need to know the specific VLANs that are used and all of the VLAN IDs used by the virtual machines that the virtual network supports.

To utilize a Hyper-V VLAN, ensure that you are using a physical network adapter that supports 802.1q VLAN tagging. By default, the virtual network adapter in Hyper-V is in untagged mode and you might need to enable the feature on a virtual network adapter.

NOTE: By using Windows PowerShell, you can determine the mode of the vNIC (**Get-VmNetworkAdapterVlan** command) and change the mode of the vNIC (**Set-VmNetworkAdapterVlan** command). See [Get-VMNetworkAdapterVlan](#) and [Set-VMNetworkAdapterVlan](#) for details on both Windows PowerShell virtual network adapter commands.


To add a VLAN for a vSRX VM virtual network adapter:

1. Open the Hyper-V Manager by selecting **Start > Administrative Tools > Hyper-V Manager**.
2. Right-click the vSRX VM and select **Settings** from the context menu.

3. From the Settings dialog box, under the Hardware section, select the network adapter connected to the external virtual network. The Network Adapter pane appears.
4. Select **Enable virtual LAN identification**, and then enter the VLAN ID you intend to use (see [Figure 18 on page 58](#)). You can change the VLAN ID to any number or leave the default. This is the VLAN

identification number that the vSRX will use for all network communication through this network adapter.

Figure 18: Enable VLAN Identification Example



5. Click **OK**, and then click **Yes** to apply networking changes.

6. If necessary, repeat Steps 3 through 5 to add VLAN identification to additional network adapters in use by the vSRX VM.

SEE ALSO

[Hyper-V: Configure VLANs and VLAN Tagging](#)

[Understanding Hyper-V VLANs](#)

Power Down a vSRX VM with Hyper-V

In situations where you need to modify the vSRX VM settings from Hyper-V, you must first perform a graceful shut down of the vSRX VM using the **Shut Down** command. The vSRX VM performs an orderly closing of all programs and attempts to shut off power to avoid data loss.

NOTE: If you are using Microsoft PowerShell, use the **Stop-VM** command to perform a graceful shutdown of the vSRX VM.

To gracefully shut down the vSRX instance on the Hyper-V host computer:

1. Log onto your Hyper-V host computer using the Administrator account.
2. Open the Hyper-V Manager by selecting **Start > Administrative Tools > Hyper-V Manager**.
3. Power down the vSRX instance in the Hyper-V Manager by selecting the vSRX VM from the list of virtual machines, and then right-click and select **Shut Down** from the context menu (or select **Action > Shut Down**).
4. Power on the vSRX instance in the Hyper-V Manager by selecting the vSRX VM from the list of virtual machines, and then right-click and select **Start** from the context menu (or select **Action > Start**).

NOTE: If you are using Microsoft PowerShell, use the **Start-VM** command to start the vSRX VM.

4

CHAPTER

Configuring and Managing vSRX

vSRX Configuration and Management Tools | 61

Configure vSRX Using the CLI | 62

Configure vSRX Using the J-Web Interface | 64

Managing Security Policies for Virtual Machines Using Junos Space Security Director | 68

Software Receive Side Scaling | 69

Multi-Core Scaling Support with SWRSS | 71

GTP Traffic with TEID Distribution and SWRSS | 72

vSRX Configuration and Management Tools

SUMMARY

This topic provides an overview of the various tools available to configure and manage a vSRX VM once it has been successfully deployed.

IN THIS SECTION

- [Understanding the Junos OS CLI and Junos Scripts | 61](#)
- [Understanding the J-Web Interface | 61](#)
- [Understanding Junos Space Security Director | 61](#)

Understanding the Junos OS CLI and Junos Scripts

Junos OS CLI is a Juniper Networks specific command shell that runs on top of a UNIX-based operating system kernel.

Built into Junos OS, Junos script automation is an onboard toolset available on all Junos OS platforms, including routers, switches, and security devices running Junos OS (such as a vSRX instance).

You can use the Junos OS CLI and the Junos OS scripts to configure, manage, administer, and troubleshoot vSRX.

Understanding the J-Web Interface

The *J-Web* interface allows you to monitor, configure, troubleshoot, and manage vSRX instances by means of a Web browser. J-Web provides access to all the configuration statements supported by the vSRX instance.

Understanding Junos Space Security Director

As one of the Junos Space Network Management Platform applications, Junos Space Security Director helps organizations improve the reach, ease, and accuracy of security policy administration with a scalable, GUI-based management tool. Security Director automates security provisioning of a vSRX

instance through one centralized Web-based interface to help administrators manage all phases of the security policy life cycle more quickly and intuitively, from policy creation to remediation.

RELATED DOCUMENTATION

[CLI User Interface Overview](#)[J-Web Overview](#)[Security Director](#)[Mastering Junos Automation Programming](#)[Spotlight Secure Threat Intelligence](#)

Configure vSRX Using the CLI

To configure the instance using the CLI:

1. Verify that the vSRX instance is powered on.
2. Log in as the root user (whose username is *root*). There is no password.
3. Start the CLI.

```
root#cli
root@>
```

4. Enter configuration mode.

```
configure
[edit]
root@#
```

5. Set the root authentication password by entering a *cleartext* password, an encrypted password, or an SSH public key string (*DSA* or *RSA*). The following is an example of a plain-text password. The CLI prompts you for the password and then encrypts it.

```
[edit]
root@# set system root-authentication plain-text-password
```

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

6. Configure the hostname.

```
[edit]
root@# set system host-name host-name
```

7. Configure the management interface.

```
[edit]
root@# set interfaces fxp0 unit 0 family inet dhcp-client
```

8. Configure the traffic interfaces.

```
[edit]
root@# set interfaces ge-0/0/0 unit 0 family inet dhcp-client
```

9. Configure basic security zones and bind them to traffic interfaces.

```
[edit]
root@# set security zones security-zone trust interfaces ge-0/0/0.0
```

10. Verify the configuration changes.

```
[edit]
root@# commit check
configuration check succeeds
```

11. Commit the configuration to activate it on the instance.

```
[edit]
root@# commit
commit complete
```

NOTE: Certain Junos OS software features require a license to activate the feature. To enable a licensed feature, you need to purchase, install, manage, and verify a license key that corresponds

to each licensed feature. To conform to software feature licensing requirements, you must purchase one license per feature per instance. The presence of the appropriate software unlocking key on your virtual instance allows you to configure and use the licensed feature. See [Managing Licenses for vSRX](#) for details.

RELATED DOCUMENTATION

[CLI User Guide](#)

[Junos OS for SRX Series](#)

Configure vSRX Using the J-Web Interface

IN THIS SECTION

- [Access the J-Web Interface and Configuring vSRX | 64](#)
- [Apply the Configuration | 67](#)
- [Add vSRX Feature Licenses | 68](#)

Access the J-Web Interface and Configuring vSRX

To configure vSRX using the *J-Web* Interface:

1. Launch the J-Web interface from a Web browser.

NOTE: You will be prompted to accept a system-generated certificate to access a vSRX VM using the J-Web interface.

2. Enter the vSRX out-of-band management (fxp0) interface IP address in the Address box.
3. Specify the username and password.

4. Click **Log In**, and select the **Configuration Wizards** tab from the left navigation panel. The J-Web Setup wizard page opens.

5. Click **Setup**.

You can use the Setup wizard to configure the vSRX VM or edit an existing configuration.

- Select **Edit Existing Configuration** if you have already configured the wizard using the factory mode.
- Select **Create New Configuration** to configure the vSRX VM using the wizard.

The following configuration options are available in the guided setup:

- Basic

Select **basic** to configure the vSRX VM name and user account information as shown in [Table 10 on page 65](#).

- Instance name and user account information

Table 10: Instance Name and User Account Information

Field	Description
Instance name	Type the name of the vSRX instance.
Root password	Create a default root user password.
Verify password	Verify the default root user password.

Table 10: Instance Name and User Account Information (*Continued*)

Field	Description
Operator	Add an optional administrative account in addition to the root account. User role options include: • Super User: This user has full system administration rights and can add, modify, and delete settings and users. • Operator: This user can perform system operations such as a system reset but cannot change the configuration or add or modify users. • Read only: This user can only access the system and view the configuration. • Disabled: This user cannot access the system.

- Select either **Time Server** or **Manual**. [Table 11 on page 66](#) lists the system time options.

Table 11: System Time Options

Field	Description
Time Server	
Host Name	Type the hostname of the time server. For example: ntp.example.com .
IP	Type the IP address of the time server in the IP address entry field. For example: 192.0.2.254 .
NOTE: You can enter either the hostname or the IP address.	
Manual	
Date	Click the current date in the calendar.

Table 11: System Time Options *(Continued)*

Field	Description
Time	Set the hour, minute, and seconds. Choose AM or PM .
Time Zone (mandatory)	
Time Zone	Select the time zone from the list. For example: GMT Greenwich Mean Time GMT.

- Expert
 - a. Select **Expert** to configure the basic options as well as the following advanced options:
 - Four or more internal zones
 - Internal zone services
 - Application of security policies between internal zones
 - b. Click the **Need Help** icon for detailed configuration information.

You see a success message after the basic configuration is complete.

Apply the Configuration

To apply the configuration settings for vSRX:

1. Review and ensure that the configuration settings are correct, and click **Next**. The Commit Configuration page appears.
2. Click **Apply Settings** to apply the configuration changes to vSRX.
3. Check the connectivity to vSRX, as you might lose connectivity if you have changed the management zone IP. Click the URL for reconnection instructions on how to reconnect to the instance.
4. Click **Done** to complete the setup.

After successful completion of the setup, you are redirected to the J-Web interface.



CAUTION: After you complete the initial setup, you can relaunch the J-Web Setup wizard by clicking **Configuration>Setup**. You can either edit an existing configuration or create a new configuration. If you create a new configuration, the current configuration in vSRX will be deleted.

Add vSRX Feature Licenses

Certain Junos OS software features require a license to activate the feature. To enable a licensed feature, you need to purchase, install, manage, and verify a license key that corresponds to each licensed feature. To conform to software feature licensing requirements, you must purchase one license per feature per instance. The presence of the appropriate software unlocking key on your virtual instance allows you to configure and use the licensed feature.

See [Managing Licenses for vSRX](#) for details.

Managing Security Policies for Virtual Machines Using Junos Space Security Director

SUMMARY

This topic provides you an overview of how you can manage security policies for VMs using security director.

Security Director is a Junos Space management application designed to enable quick, consistent, and accurate creation, maintenance, and application of network security policies for your security devices, including vSRX instances. With Security Director, you can configure security-related policy management including IPsec VPNs, firewall policies, NAT policies, IPS policies, and UTM policies. and push the configurations to your security devices. These configurations use objects such as addresses, services, NAT pools, application signatures, policy profiles, VPN profiles, template definitions, and templates. These objects can be shared across multiple security configurations; shared objects can be created and

used across many security policies and devices. You can create these objects prior to creating security configurations.

When you finish creating and verifying your security configurations from Security Director, you can publish these configurations and keep them ready to be pushed to all security devices, including vSRX instances, from a single interface.

The Configure tab is the workspace where all of the security configuration happens. You can configure firewall, IPS, NAT, and UTM policies; assign policies to devices; create and apply policy schedules; create and manage VPNs; and create and manage all the shared objects needed for managing your network security.

RELATED DOCUMENTATION

[Security Director](#)

Software Receive Side Scaling

IN THIS SECTION

- [Overview](#) | 69
- [Understanding Software Receive Side Scaling Configuration](#) | 70

Overview

Contemporary NICs support multiple receive and transmit descriptor queues (multi-queue). On reception, a NIC can send different packets to different queues to distribute processing among CPUs. The NIC distributes packets by applying a filter to each packet that assigns it to one of a small number of logical flows. Packets for each flow are steered to a separate receive queue, which in turn can be processed by separate CPUs. This mechanism is generally known as Receive-side Scaling (RSS). The goal of RSS technique is to increase performance uniformly. RSS is enabled when latency is a concern or whenever receive interrupt processing forms a bottleneck. Spreading load between CPUs decreases queue length. For low latency networking, the optimal setting is to allocate as many queues as there are CPUs in the system (or the NIC maximum, if lower). The most efficient high-rate configuration is likely

the one with the smallest number of receive queues where no receive queue overflows due to a saturated CPU. You can improve bridging throughput with Receive Side Scaling.

As per flow thread affinity architecture each flow thread (FLT) polls for packet from dedicated receiving queue of NIC and process the packets until run to completion. Therefore, flow threads are bound to NIC receiving (RX) and transmitting (TX) queues for packet processing to avoid any disagreement. Hence, NIC must have same number of RX and TX queues as number of vSRX data plane CPU to support multi-core vSRX flavors. Software RSS (SWRSS) removes this limitation of NIC HW queues to run vSRX multi-core flavors by implementing software-based packet spraying across various FLT thread.

Software RSS offloads the handling of individual flows to one of the multiple kernel, so the flow thread that takes the packets from the NIC can process more packets. Similar to RSS, network throughput improvement when using SWRSS has a linear correlation with CPU utilization.

In SWRSS, each NIC port is initialized with equal or lesser number of hardware RX/TX queues as that of I/O threads. I/O threads are determined based on total data-path CPU and minimum of NIC queues among all the NIC interface in vSRX. For example, if I/O thread is computed as 4, then number of HW queue per NIC port can be less or equal to 4 queues.

If NICs do not have sufficient number of queues as FLT threads in vSRX instances supported, then Software RSS (SWRSS) is enabled by flowd data-path. SWRSS implements software model of packet distribution across FLTs after obtaining the packets from NIC receiving queues. By removing NIC HW queue limitation, SWRSS helps to scale vCPUs by supporting various vSRX instance types.

During the I/O operation the packets are fetched from receiving queues of NIC ports and packet classification is performed. Followed by distribution of packets to FLT threads virtual queues. These virtual queues are implemented over DPDK ring queue. In the transmission path, SWRSS fetches the packets from virtual transmitting queues of FLT threads and pushes these packets to NIC transmitting queues for transmit.

Number of SWRSS I/O threads are selected based on total CPU and number of NIC queues found in vSRX instances. Mix mode of operation with HWRSS and SWRSS is not supported.

Understanding Software Receive Side Scaling Configuration

This topic provide you details on types of Software Receive Side Scaling (SWRSS) and its configuration.

SWRSS supports two modes of operation and it gets enabled based on number of data-path CPU needed. These modes are Shared IO mode and dedicated IO mode. These modes are enabled based on number of data-path CPUs needed. vSRX and vSRX3.0 supports dedicated I/O mode only.

In dedicated I/O mode flowd process creates dedicated I/O threads for I/O operation. Based on number of required I/O threads for vSRX, I/O thread is associated to a dedicated NIC port. NIC ports receiving and transmitting queue is then bonded to each I/O thread in round robin method for uniform

distribution and to avoid I/O thread locks. Each dedicated I/O thread pulls the packets in burst mode from NIC receiving queue and distributes to FLT threads and vice versa for TX path for packet transmit.

SWRSS is enabled based on the number of vCPUs. If NIC does not have sufficient number of queues as flow thread (FLT) in vSRX with different vCPUs, then Software RSS (SWRSS) is enabled by flowd process.

SWRSS is not enabled in the following scenarios:

- When the NIC has sufficient number of hardware RX or TX queues for required PFE data-path CPU.
- When the vSRX (based on number of vCPUs) and NIC result the smaller number of FLT CPUs as that obtained in nearest hardware RSS (HWRSS) mode. In such scenario, vSRX will be enabled with HWRSS mode which results more FLT CPU than SWRSS mode, providing better packet processing throughput.
- SWRSS is not recommended for vSRX with certain type of NIC that supports lesser number of NIC queues than needed to run dedicated IO thread. In such cases, SWRSS is enabled but extra CPUs are attached to FLT CPU, until I/O CPUs are completely utilized.

If SWRSS is not enabled use the **set security forwarding-options receive-side-scaling software-rss mode enable** command to enable SWRSS. When you run this command SWRSS will be enabled by force regardless of the NIC RSS or the number of vCPUs. If you do not enable SWRSS using the CLI then enabling of SWRSS automatically is decided based on the default ratio of FLT: IO (4:1).

To configure the number of required IO threads, use the **set security forwarding-options receive-side-scaling software-rss io-thread-number <1-8>** command. To view the actual number of vCPUs assigned to IO flow threads use the **show security forwarding-options resource-manager** command.

You can decide enabling of SWRSS automatically or by force based on the architecture and conception of IO thread and worker thread. Enabling SWRSS impacts the performance, so we recommend that the number of IO thread should be changed only if required and until the performance impact bottleneck point is reached.

Multi-Core Scaling Support with SWRSS

IN THIS SECTION

- [Limitations | 72](#)

Starting in Junos OS Release 19.4R1, vSRX 3.0 instances on Microsoft Hyper-V with the Software Receive Side Scaling (SWRSS) feature can scale up based on the vCPUs supported.

With SWRSS, the vCPUs are specified as IO thread usage and worker thread usage. The vSRX3.0 instances running on Hyper-V use a dedicated IO mode of SWRSS.

Limitations

For SWRSS, up to 32 vCPUs can be used for each vSRX instance, and the ratio of worker thread and the IO thread is 4:1 to deliver better performance.

GTP Traffic with TEID Distribution and SWRSS

IN THIS SECTION

- [Overview GTP Traffic Distribution with TEID Distribution and SWRSS | 72](#)
- [Enabling GTP-U TEID Distribution with SWRSS for Asymmetric Fat Tunnels | 74](#)

Overview GTP Traffic Distribution with TEID Distribution and SWRSS

IN THIS SECTION

- [GTP Traffic Performance with TEID Distribution and SWRSS | 73](#)

The topic provides an overview of asymmetric fat tunnel solution for GTP traffic with TEID distribution and SWRSS.

With TEID-based hash distributions feature, the GTP packets would be distributed to the flow thread according to the hash value calculated by TEID. The algorithm of hash calculation is same as GTP

distribution in flow module, which ensures the GTP packets would not be reinjected again in the flow process.

There is a 4-byte field inside GTP payload called tunnel endpoint identifier (TEID), which is used to identify different connections in the same GTP tunnel.

A fat GTP tunnel carries data from different users. IPsec tunnels on the security gateway could be a fat tunnel due to the fat GTP tunnel. vSRX can create one GTP session with a high-bandwidth of GTP traffic. However, the throughput is limited to one core processor's performance.

If you use TEID-based hash distribution for creating GTP-U sessions, then you can:

- Enable vSRX and vSRX 3.0 instances to process asymmetric fat tunnels for parallel encryption on multiple cores for one tunnel.
- You can split a fat GTP session to multiple sessions and distribute them to different cores. This helps to increase the bandwidth for fat GTP tunnel.

The TEID based hash distribution creates GTP-U sessions to multiple cores. The clear text traffic acts as a fat GTP tunnel. This helps a fat GTP session to split into multiple slim GTP sessions and handle them on multiple cores simultaneously.

GTP Traffic Performance with TEID Distribution and SWRSS

vSRX instances support Software Receive Side Scaling (SWRSS) feature. SWRSS is a technique in the networking stack to increase parallelism and improve performance for multi-processor systems. If NICs do not have sufficient number of queues as flow thread (FLT), based on vSRX type, then Software RSS (SWRSS) is enabled by flowd process.

With Software Receive Side Scaling (SWRSS) support on vSRX and vSRX 3.0, you can assign more vCPUs to the vSRX regardless of the limitation of RSS queue of underlying interfaces.

Based on SWRSS you can improve the GTP traffic performance using Tunnel endpoint identifier (TEID) distribution and asymmetric fat tunnel solution by:

- Assigning specific number of vCPUs for input output flow usage—With SWRSS enabled, you can assign more vCPUs for input/output (IO) threads when the IO threads are less. Or you can assign less vCPUs for IO threads if the flow process is consuming more vCPU. Use the **set security forwarding-options receive-side-scaling software-rss io-thread-number <io-thread-number>**.
- Distributing the packets to flow threads according to the TEID inside the packet, which would avoid reinjecting the packets in flow process—This feature is enabled when both SWRSS is enabled and when you configure the **set security forwarding-process application-services enable-gtpu-distribution** command.

With this feature, the GTP packets would be distributed to the flow thread according to the hash value calculated by TEID. The algorithm of hash calculation is same as GTP distribution in flow module, which ensures the GTP packets would not be reinjected again in flow process.

- Utilizing fragment matching and forwarding mechanism in input/output thread when GTPU distribution is enabled—This mechanism ensures that all the fragments of the same packet would be distributed to one flow thread according to the TEID.

SWRSS uses IP pair hash to distribute packets to flow threads. For GTP traffic with GTPU distribution enabled, TEID distribution is used to distribute packets to the flow threads. For fragmented packets, TEID cannot be retrieved from non-first fragments. This will require fragment matching and forwarding logic to ensure all fragments are forwarded to the flow thread based on TEID.

Enabling GTP-U TEID Distribution with SWRSS for Asymmetric Fat Tunnels

The following configuration helps you enable PMI and GTP-U traffic distribution with SWRSS enabled.

Before you begin, understand:

- SWRSS concepts and configurations.
- How to establish PMI and GTP-U

With Software Recieve Side Scaling (SWRSS) enabled, you can assign more vCPUs for input/output (IO) threads when the IO threads are less. Or you can assign less vCPUs for IO threads if the flow process is consuming more vCPU. You can configure the number of IO threads required. With SWRSS is enabled and IO threads configured, reboot the vSRX for configuration to take effect. After IO threads are configured, distribute the GTP traffic to the configured IO threads according to TEID-based hash distribution for splitting a fat GTP session to multiple slim GTP sessions and process them on multiple cores in parallel.

NOTE: When PMI mode is enabled with TEID distribution and SWRSS support, performance of PMI is improved. If you want to enable PMI mode then run the **set securtiy flow power-mode-ipsec** command.

The following steps provide you details on how to enable SWRSS, configure IO threads, enable PMI mode for GTP sessions with TEID distribution for obtaining asymmetric fat tunnels:

1. SWRSS is enabled by default when NICs do not have sufficient number of queues as flow thread (FLT) based on vSRX type, then Software RSS (SWRSS) is enabled by flowd process. But, when SWRSS is not enabled use the following CLIs to enable. When you run this command SWRSS will be enabled by force regardless of the NIC RSS or number of vCPUs.

Enable SWRSS.

[edit]

```
user@host# set security forwarding-options receive-side-scaling software-rss mode enable
```

2. Configure the number of IO threads required. In this configuration we are configuring eight IO threads. The assigned number of vCPUs would be assigned for IO threads, and the rest vCPUs would be assigned for flow thread.

[edit]

```
user@host# set security forwarding-options receive-side-scaling software-rss io-thread-number 8
```

- 3.

[edit security]

```
user@host# set flow power-mode-ipsec
```

4. Configure GTP-U session distribution.

[edit security]

```
user@host# set forwarding-process application-services enable-gtpu-distribution
```

5. From the configuration mode, confirm your configuration by entering the **show** command.

[edit security]

```
user@host# show
forwarding-options {
  receive-side-scaling {
    software-rss {
      mode enable;
      io-thread-number 8;
    }
  }
  flow {
    power-mode-ipsec;
  }
  forwarding-process {
    application-services {
```

```

        enable-gtpu-distribution;
    }
}

```

From the operational mode run the following command to view the actual number of vCPUs assigned to IO/flow threads.

```
show security forward-options resource-manager settings
```

```

-----
Owner          Type                Current settings  Next settings
SWRSS-IO       CPU core number    2                 2
SWRSS          SWRSS mode         Enable            Enable

```

6. Commit the configuration.

```

[edit security]
user@host# commit

```

7. Reboot the vSRX for the configuration to take effect. After rebooting the whole device, PFE would check the IO-thread value according to the NIC RSS queue and its memory.

5

CHAPTER

Configuring vSRX Chassis Clusters

[Configure a vSRX Chassis Cluster in Junos OS | 78](#)

[vSRX Cluster Staging and Provisioning in Hyper-V | 88](#)

Configure a vSRX Chassis Cluster in Junos OS

IN THIS SECTION

- [Chassis Cluster Overview | 78](#)
- [Enable Chassis Cluster Formation | 79](#)
- [Chassis Cluster Quick Setup with J-Web | 80](#)
- [Manually Configure a Chassis Cluster with J-Web | 81](#)

Chassis Cluster Overview

Prerequisites

Ensure that your vSRX instances comply with the following prerequisites before you enable chassis clustering:

- Use **show version** in Junos OS to ensure that both vSRX instances have the same software version.
- Use **show system license** in Junos OS to ensure that both vSRX instances have the same licenses installed.

Chassis cluster groups a pair of the same kind of vSRX instances into a cluster to provide network node redundancy. The devices must be running the same Junos OS release. You connect the control virtual interfaces on the respective nodes to form a *control plane* that synchronizes the configuration and Junos OS kernel state. The control link (a *virtual network* or *vSwitch*) facilitates the redundancy of interfaces and services. Similarly, you connect the *data plane* on the respective nodes over the fabric virtual interfaces to form a unified data plane. The fabric link (a virtual network or vSwitch) allows for the management of cross-node flow processing and for the management of session redundancy.

The control plane software operates in active/passive mode. When configured as a chassis cluster, one node acts as the primary device and the other as the secondary device to ensure stateful failover of processes and services in the event of a system or hardware failure on the primary device. If the primary device fails, the secondary device takes over processing of control plane traffic.

NOTE: If you configure a chassis cluster on vSRX nodes across two physical hosts, disable igmp-snooping on the bridge that each host physical interface belongs to that the control vNICs use. This ensures that the control link heartbeat is received by both nodes in the chassis cluster.

The chassis cluster data plane operates in active/active mode. In a chassis cluster, the data plane updates session information as traffic traverses either device, and it transmits information between the nodes over the fabric link to guarantee that established sessions are not dropped when a failover occurs. In active/active mode, traffic can enter the cluster on one node and exit from the other node.

Chassis cluster functionality includes:

- Resilient system architecture, with a single active control plane for the entire cluster and multiple *Packet Forwarding Engines*. This architecture presents a single device view of the cluster.
- Synchronization of configuration and dynamic runtime states between nodes within a cluster.
- Monitoring of physical interfaces, and failover if the failure parameters cross a configured threshold.
- Support for generic routing encapsulation (*GRE*) and IP-over-IP (IP-IP) tunnels used to route encapsulated IPv4 or *IPv6* traffic by means of two internal interfaces, *gr-0/0/0* and *ip-0/0/0*, respectively. Junos OS creates these interfaces at system startup and uses these interfaces only for processing GRE and IP-IP tunnels.

At any given instant, a cluster node can be in one of the following states: hold, primary, secondary-hold, secondary, ineligible, or disabled. Multiple event types, such as interface monitoring, Services Processing Unit (SPU) monitoring, failures, and manual failovers, can trigger a state transition.

Enable Chassis Cluster Formation

You create two vSRX instances to form a chassis cluster, and then you set the cluster ID and node ID on each instance to join the cluster. When a vSRX VM joins a cluster, it becomes a node of that cluster. With the exception of unique node settings and management IP addresses, nodes in a cluster share the same configuration.

You can deploy up to 255 chassis clusters in a *Layer 2* domain. Clusters and nodes are identified in the following ways:

- The *cluster ID* (a number from 1 to 255) identifies the cluster.
- The *node ID* (a number from 0 to 1) identifies the cluster node.

On SRX Series devices, the cluster ID and node ID are written into EEPROM. On the vSRX VM, vSRX stores and reads the IDs from **boot/loader.conf** and uses the IDs to initialize the chassis cluster during startup.

The chassis cluster formation commands for node 0 and node 1 are as follows:

- On vSRX node 0:

```
user@vsrx0>set chassis cluster cluster-id number node 0 reboot
```

- On vSRX node 1:

```
user@vsrx1>set chassis cluster cluster-id number node 1 reboot
```

NOTE: Use the same cluster ID number for each node in the cluster.

NOTE: The vSRX interface naming and mapping to vNICs changes when you enable chassis clustering.

After reboot, on node 0, configure the fabric (data) ports of the cluster that are used to pass real-time objects (RTOs):

- ```
user@vsrx0# set interfaces fab0 fabric-options member-interfaces ge-0/0/0
user@vsrx0# set interfaces fab1 fabric-options member-interfaces ge-7/0/0
```

## Chassis Cluster Quick Setup with J-Web

To configure chassis cluster from *J-Web*:

1. Enter the vSRX node 0 interface IP address in a Web browser.
2. Enter the vSRX username and password, and click **Log In**. The J-Web dashboard appears.
3. Click **Configuration Wizards>Chassis Cluster** from the left panel. The Chassis Cluster Setup wizard appears. Follow the steps in the setup wizard to configure the cluster ID and the two nodes in the cluster, and to verify connectivity.

**NOTE:** Use the built-in Help icon in J-Web for further details on the Chassis Cluster Setup wizard.

## Manually Configure a Chassis Cluster with J-Web

You can use the *J-Web* interface to configure the primary node 0 vSRX instance in the cluster. Once you have set the cluster and node IDs and rebooted each vSRX, the following configuration will automatically be synced to the secondary node 1 vSRX instance.

Select **Configure>Chassis Cluster>Cluster Configuration**. The Chassis Cluster configuration page appears.

[Table 12 on page 81](#) explains the contents of the HA Cluster Settings tab.

[Table 13 on page 83](#) explains how to edit the Node Settings tab.

[Table 14 on page 84](#) explains how to add or edit the HA Cluster Interfaces table.

[Table 15 on page 85](#) explains how to add or edit the HA Cluster Redundancy Groups table.

**Table 12: Chassis Cluster Configuration Page**

| Field                | Function                                                                                                                          |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------|
| <b>Node Settings</b> |                                                                                                                                   |
| Node ID              | Displays the node ID.                                                                                                             |
| Cluster ID           | Displays the cluster ID configured for the node.                                                                                  |
| Host Name            | Displays the name of the node.                                                                                                    |
| Backup Router        | Displays the router used as a gateway while the Routing Engine is in secondary state for redundancy-group 0 in a chassis cluster. |
| Management Interface | Displays the management interface of the node.                                                                                    |

Table 12: Chassis Cluster Configuration Page *(Continued)*

| Field      | Function                                                                                                                                                                                            |
|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IP Address | Displays the management IP address of the node.                                                                                                                                                     |
| Status     | Displays the state of the redundancy group. <ul style="list-style-type: none"> <li>• <b>Primary</b>—Redundancy group is active.</li> <li>• <b>Secondary</b>—Redundancy group is passive.</li> </ul> |

## Chassis Cluster&gt;HA Cluster Settings&gt;Interfaces

|                              |                                                                               |
|------------------------------|-------------------------------------------------------------------------------|
| Name                         | Displays the physical interface name.                                         |
| Member Interfaces/IP Address | Displays the member interface name or IP address configured for an interface. |
| Redundancy Group             | Displays the redundancy group.                                                |

## Chassis Cluster&gt;HA Cluster Settings&gt;Redundancy Group

|                      |                                                                                                                                                                                                                                    |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Group                | Displays the redundancy group identification number.                                                                                                                                                                               |
| Preempt              | Displays the selected preempt option. <ul style="list-style-type: none"> <li>• <b>True</b>—Primary Role can be preempted based on priority.</li> <li>• <b>False</b>—Primary Role cannot be preempted based on priority.</li> </ul> |
| Gratuitous ARP Count | Displays the number of gratuitous Address Resolution Protocol (ARP) requests that a newly elected primary device in a chassis cluster sends out to announce its presence to the other network devices.                             |

Table 12: Chassis Cluster Configuration Page (*Continued*)

| Field         | Function                                                                                                                                                         |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Node Priority | Displays the assigned priority for the redundancy group on that node. The eligible node with the highest priority is elected as primary for the redundant group. |

Table 13: Edit Node Setting Configuration Details

| Field                | Function                                                                                                                              | Action                                     |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------|
| <b>Node Settings</b> |                                                                                                                                       |                                            |
| Host Name            | Specifies the name of the host.                                                                                                       | Enter the name of the host.                |
| Backup Router        | Displays the device used as a gateway while the Routing Engine is in the secondary state for redundancy-group 0 in a chassis cluster. | Enter the IP address of the backup router. |
| <b>Destination</b>   |                                                                                                                                       |                                            |
| IP                   | Adds the destination address.                                                                                                         | Click <b>Add</b> .                         |
| Delete               | Deletes the destination address.                                                                                                      | Click <b>Delete</b> .                      |
| <b>Interface</b>     |                                                                                                                                       |                                            |
| Interface            | Specifies the interfaces available for the router.<br><b>NOTE:</b> Allows you to add and edit two interfaces for each fabric link.    | Select an option.                          |
| IP                   | Specifies the interface IP address.                                                                                                   | Enter the interface IP address.            |

Table 13: Edit Node Setting Configuration Details *(Continued)*

| Field  | Function               | Action                |
|--------|------------------------|-----------------------|
| Add    | Adds the interface.    | Click <b>Add</b> .    |
| Delete | Deletes the interface. | Click <b>Delete</b> . |

Table 14: Add HA Cluster Interface Configuration Details

| Field | Function | Action |
|-------|----------|--------|
|-------|----------|--------|

**Fabric Link > Fabric Link 0 (fab0)**

|           |                             |                                       |
|-----------|-----------------------------|---------------------------------------|
| Interface | Specifies fabric link 0.    | Enter the interface IP fabric link 0. |
| Add       | Adds fabric interface 0.    | Click <b>Add</b> .                    |
| Delete    | Deletes fabric interface 0. | Click <b>Delete</b> .                 |

**Fabric Link > Fabric Link 1 (fab1)**

|           |                             |                                           |
|-----------|-----------------------------|-------------------------------------------|
| Interface | Specifies fabric link 1.    | Enter the interface IP for fabric link 1. |
| Add       | Adds fabric interface 1.    | Click <b>Add</b> .                        |
| Delete    | Deletes fabric interface 1. | Click <b>Delete</b> .                     |

**Redundant Ethernet**

|           |                                                                                                    |                              |
|-----------|----------------------------------------------------------------------------------------------------|------------------------------|
| Interface | Specifies a logical interface consisting of two physical Ethernet interfaces, one on each chassis. | Enter the logical interface. |
|-----------|----------------------------------------------------------------------------------------------------|------------------------------|

Table 14: Add HA Cluster Interface Configuration Details *(Continued)*

| Field            | Function                                                         | Action                                   |
|------------------|------------------------------------------------------------------|------------------------------------------|
| IP               | Specifies a redundant Ethernet IP address.                       | Enter a redundant Ethernet IP address.   |
| Redundancy Group | Specifies the redundancy group ID number in the chassis cluster. | Select a redundancy group from the list. |
| Add              | Adds a redundant Ethernet IP address.                            | Click <b>Add</b> .                       |
| Delete           | Deletes a redundant Ethernet IP address.                         | Click <b>Delete</b> .                    |

Table 15: Add Redundancy Groups Configuration Details

| Field                           | Function                                                                                                                                                                                                                                                                                                                                                                | Action                                        |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------|
| Redundancy Group                | Specifies the redundancy group name.                                                                                                                                                                                                                                                                                                                                    | Enter the redundancy group name.              |
| Allow preemption of primaryship | <p>Allows a node with a better priority to initiate a failover for a redundancy group.</p> <p><b>NOTE:</b> By default, this feature is disabled. When disabled, a node with a better priority does not initiate a redundancy group failover (unless some other factor, such as faulty network connectivity identified for monitored interfaces, causes a failover).</p> | –                                             |
| Gratuitous ARP Count            | Specifies the number of gratuitous Address Resolution Protocol requests that a newly elected primary sends out on the active redundant Ethernet interface child links to notify network devices of a change in primary role on the redundant Ethernet interface links.                                                                                                  | Enter a value from 1 to 16. The default is 4. |

Table 15: Add Redundancy Groups Configuration Details (*Continued*)

| Field                    | Function                                                                                        | Action                                                                  |
|--------------------------|-------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------|
| node0 priority           | Specifies the priority value of node0 for a redundancy group.                                   | Enter the node priority number as 0.                                    |
| node1 priority           | Specifies the priority value of node1 for a redundancy group.                                   | Select the node priority number as 1.                                   |
| <b>Interface Monitor</b> |                                                                                                 |                                                                         |
| Interface                | Specifies the number of redundant Ethernet interfaces to be created for the cluster.            | Select an interface from the list.                                      |
| Weight                   | Specifies the weight for the interface to be monitored.                                         | Enter a value from 1 to 125.                                            |
| Add                      | Adds interfaces to be monitored by the redundancy group along with their respective weights.    | Click <b>Add</b> .                                                      |
| Delete                   | Deletes interfaces to be monitored by the redundancy group along with their respective weights. | Select the interface from the configured list and click <b>Delete</b> . |

**IP Monitoring**

|             |                                                                         |                              |
|-------------|-------------------------------------------------------------------------|------------------------------|
| Weight      | Specifies the global weight for IP monitoring.                          | Enter a value from 0 to 255. |
| Threshold   | Specifies the global threshold for IP monitoring.                       | Enter a value from 0 to 255. |
| Retry Count | Specifies the number of retries needed to declare reachability failure. | Enter a value from 5 to 15.  |

Table 15: Add Redundancy Groups Configuration Details (*Continued*)

| Field                                 | Function                                                                  | Action                                                          |
|---------------------------------------|---------------------------------------------------------------------------|-----------------------------------------------------------------|
| Retry Interval                        | Specifies the time interval in seconds between retries.                   | Enter a value from 1 to 30.                                     |
| <b>IPv4 Addresses to Be Monitored</b> |                                                                           |                                                                 |
| IP                                    | Specifies the IPv4 addresses to be monitored for reachability.            | Enter the IPv4 addresses.                                       |
| Weight                                | Specifies the weight for the redundancy group interface to be monitored.  | Enter the weight.                                               |
| Interface                             | Specifies the logical interface through which to monitor this IP address. | Enter the logical interface address.                            |
| Secondary IP address                  | Specifies the source address for monitoring packets on a secondary link.  | Enter the secondary IP address.                                 |
| Add                                   | Adds the IPv4 address to be monitored.                                    | Click <b>Add</b> .                                              |
| Delete                                | Deletes the IPv4 address to be monitored.                                 | Select the IPv4 address from the list and click <b>Delete</b> . |

**SEE ALSO**

[Chassis Cluster Feature Guide for Security Devices](#)

# vSRX Cluster Staging and Provisioning in Hyper-V

## IN THIS SECTION

- [Deploying the VMs and Additional Network Adapters in Hyper-V | 88](#)
- [Creating the Control Link Connection in Hyper-V | 89](#)
- [Creating the Fabric Link Connection in Hyper-V | 93](#)
- [Creating the Data Interfaces Using Hyper-V | 93](#)
- [Prestaging the Configuration from the Console | 94](#)
- [Connecting and Installing the Staging Configuration | 96](#)

Staging and provisioning a vSRX cluster on a Hyper-V host computer includes the following tasks:

**NOTE:** Starting in Junos OS Release 15.1X49-D100 and Junos OS Release 17.4R1, support for chassis clustering to provide network node redundancy is only available on Windows Hyper-V Server 2016.

## Deploying the VMs and Additional Network Adapters in Hyper-V

The vSRX cluster uses three interfaces exclusively for clustering (the first two are predefined):

- Out-of-band management interface (fxp0).
- Cluster control link (em0).
- Cluster fabric links (fab0 and fab1). For example, you can specify ge-0/0/0 as fab0 on node0 and ge-7/0/0 as fab1 on node1.

A cluster requires three interfaces (two for the cluster and one for management) and additional interfaces to forward data. This section outlines how to create the control link and fabric link connections, and to map all data interfaces to network adapters.

**NOTE:** For an overview on the procedure to add virtual switches and map the virtual switch to a network adapter, see ["Add vSRX Interfaces" on page 49](#)


## Creating the Control Link Connection in Hyper-V

To connect the control interface through the control link virtual switch using Hyper-V Manager:

1. Open the Hyper-V Manager by selecting **Start > Administrative Tools > Hyper-V Manager**.
2. From the Hyper-V Manager, select **Action > Virtual Switch Manager**. The Virtual Switch Manager appears.

3. Under the Virtual Switches section, select **New virtual network switch**. The Create Virtual Switch pane appears (see [Figure 19 on page 90](#)).

Figure 19: Create Virtual Switch Pane



4. Select **Internal** as the type of virtual switch. Internal allows communication between virtual machines on the same Hyper-V server, and between the virtual machines and the management host operating system.

5. Select **Create Virtual Switch**. The Virtual Switch Properties page appears (see [Figure 20 on page 91](#)).

Figure 20: Virtual Switch Properties Pane

**Virtual Switch Properties**

Name:  
ctrl\_link

Notes:

**Connection type**  
What do you want to connect this virtual switch to?

☐ External network:  
Broadcom NetXtreme Gigabit Ethernet  
☒ Allow management operating system to share this network adapter  
☐ Enable single-root I/O virtualization (SR-IOV)

☒ Internal network  
☐ Private network

**VLAN ID**  
☐ Enable virtual LAN identification for management operating system  
The VLAN identifier specifies the virtual LAN that the management operating system will use for all network communications through this network adapter. This setting does not affect virtual machine networking.  
2

Remove


**i** SR-IOV can only be configured when the virtual switch is created. An external virtual switch with SR-IOV enabled cannot be converted to an internal or private switch.

OK Cancel Apply

6. Specify a name for the control link virtual switch. Leave the other virtual switch properties at their default settings.
7. Click **OK** and then click **Yes** to apply networking changes and to close the Virtual Switch Manager window.

8. Right-click the vSRX VM and select **Settings** from the context menu. From the Settings dialog for the vSRX VM, the Hardware section, click **Network Adapter**. The Network Adapter pane appears (see [Figure 21 on page 92](#)). Assign network adapter 2 as the control link (em0) virtual switch.

Figure 21: Adding Virtual Switch to Network Adapter Pane Example



9. From the Virtual switch drop-down assign **ctrl\_link** to the control link virtual switch.
10. From the Network Adapter pane, select **Advanced Features**. Select the **Enable MAC address spoofing** check box to enable the MAC address spoofing function for the network adapter. MAC address spoofing is a requirement for the control link interface included in the redundancy groups.
11. Click **OK** and then click **Yes** to apply network adapter changes.

## Creating the Fabric Link Connection in Hyper-V

To connect the fabric interface through the fabric link virtual switch using Hyper-V Manager

1. If necessary, open the Hyper-V Manager by selecting **Start > Administrative Tools > Hyper-V Manager**.
2. From the Hyper-V Manager, select **Action > Virtual Switch Manager**. The Virtual Switch Manager appears.
3. Under the Virtual Switches section, select **New virtual network switch**. The Create Virtual Switch pane appears (see [Figure 19 on page 90](#)).
4. Select **Internal** as the type of virtual switch. Internal allows communication between virtual machines on the same Hyper-V server, and between the virtual machines and the management host operating system.
5. Select **Create Virtual Switch**. The Virtual Switch Properties page appears (see [Figure 20 on page 91](#)).
6. Specify a name for the fabric link virtual switch. Leave the other virtual switch properties at their default settings.
7. Click **OK** and then click **Yes** to apply networking changes and to close the Virtual Switch Manager window.
8. Right-click the vSRX VM and select **Settings** from the context menu. From the Settings dialog for the vSRX VM, the Hardware section, click **Network Adapter** to access the Network Adapter pane. The Network Adapter pane appears (see [Figure 21 on page 92](#)). Assign network adapter 3 as the fabric link (fab 0 or fab 1) virtual switch.
9. From the Virtual switch drop-down assign **fab0 or fab1** to the fabric link virtual switch.
10. From the Network Adapter pane, select **Advanced Features**. Select the **Enable MAC address spoofing** check box to enable the MAC address spoofing function for the network adapter. MAC address spoofing is a requirement for the fabric link interface included in the redundancy groups.
11. Click **OK** and then click **Yes** to apply network adapter changes.

## Creating the Data Interfaces Using Hyper-V

To map all data interfaces to the desired network adapters:

1. If necessary, open the Hyper-V Manager by selecting **Start > Administrative Tools > Hyper-V Manager**.
2. From the Hyper-V Manager, select **Action > Virtual Switch Manager**. The Virtual Switch Manager appears.

3. Under the Virtual Switches section, select **New virtual network switch**. The Create Virtual Switch pane appears (see [Figure 19 on page 90](#)).
4. Select **Internal** as the type of virtual switch. Internal allows communication between virtual machines on the same Hyper-V server, and between the virtual machines and the management host operating system.
5. Select **Create Virtual Switch**. The Virtual Switch Properties page appears (see [Figure 20 on page 91](#)).
6. Specify a name for the data interface virtual switch. Leave the other virtual switch properties at their default settings.
7. Click **OK** and then click **Yes** to apply networking changes and to close the Virtual Switch Manager window.
8. Right-click the vSRX VM and select **Settings** from the context menu. From the Settings dialog for the vSRX VM, the Hardware section, click **Network Adapter** to access the Network Adapter pane. The Network Adapter pane appears (see [Figure 21 on page 92](#)). Assign network adapter 3 as the data interface (fab 0 or fab 1) virtual switch.
9. From the Virtual switch drop-down assign **data interface** to the virtual switch.
10. From the Network Adapter pane, select **Advanced Features**. Select the **Enable MAC address spoofing** check box to enable the MAC address spoofing function for the network adapter. MAC address spoofing is a requirement for the data interfaces included in the redundancy groups.
11. Click **OK** and then click **Yes** to apply network adapter changes. The data interface will be connected through the data virtual switch.

## Prestaging the Configuration from the Console

The following procedure explains the configuration commands required to set up the vSRX chassis cluster. The procedure powers up both nodes, adds the configuration to the cluster, and allows SSH remote access.

1. Log in as the root user. There is no password.
2. Start the CLI.

```
root#cli
root@>
```

3. Enter configuration mode.

```
configure
[edit]
root@#
```

4. Copy the following commands and paste them into the CLI:

```
set groups node0 interfaces fxp0 unit 0 family inet address 192.168.42.81/24
set groups node0 system hostname vsrx-node0
set groups node1 interfaces fxp0 unit 0 family inet address 192.168.42.82/24
set groups node1 system hostname vsrx-node1
set apply-groups "${node}"
```

5. Set the root authentication password by entering a cleartext password, an encrypted password, or an SSH public key string (DSA or RSA).

```
root@# set system root-authentication plain-text-password
New password: password
Retype new password: password
set system root-authentication encrypted-password "$ABC123"
```

6. To enable SSH remote access:

```
user@host#set system services ssh
```

7. To enable IPv6:

```
user@host#set security forwarding-options family inet6 mode flow-based
```

This step is optional and requires a system reboot.

8. Commit the configuration to activate it on the device.

```
user@host#commit
commit complete
```

9. When you have finished configuring the device, exit configuration mode.

```
user@host#exit
```

## Connecting and Installing the Staging Configuration

After the vSRX cluster initial setup, set the cluster ID and the node ID, as described in ["Configure a vSRX Chassis Cluster in Junos OS" on page 78](#).

After reboot, the two nodes are reachable on interface fxp0 with SSH. If the configuration is operational, the **show chassis cluster status** command displays output similar to that shown in the following sample output.

vsrx> **show chassis cluster status**

```
Cluster ID: 1
Node Priority Status Preempt Manual failover

Redundancy group: 0 , Failover count: 1
 node0 100 secondary no no
 node1 150 primary no no

Redundancy group: 1 , Failover count: 1
 node0 100 secondary no no
 node1 150 primary no no
```

A cluster is healthy when the primary and secondary nodes are present and both have a priority greater than 0.

### Release History Table

| Release      | Description                                                                                                                                                                                |
|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 15.1X49-D100 | Starting in Junos OS Release 15.1X49-D100 and Junos OS Release 17.4R1, support for chassis clustering to provide network node redundancy is only available on Windows Hyper-V Server 2016. |

# 6

CHAPTER

## Troubleshooting

---

Finding the Software Serial Number for vSRX | 98

---

# Finding the Software Serial Number for vSRX

You need the software serial number to open a support case or to renew a vSRX license.

The serial number is a unique 14-digit number that Juniper Networks uses to identify your particular software installation. You can find the software serial number in the Software Serial Number Certificate attached to the e-mail that was sent when you ordered your Juniper Networks software or license. You can also use the `show system license` command to find the software serial number.

Use the **show system license** command to find the vSRX software serial number.

```
vsrx> show system license
```

License usage:

|                                   | Licenses<br>used | Licenses<br>installed | Licenses<br>needed | Expiry  |
|-----------------------------------|------------------|-----------------------|--------------------|---------|
| Feature name<br>Virtual Appliance | 1                | 1                     | 0                  | 58 days |

Licenses installed:

License identifier: E420588955

License version: 4

Software Serial Number: 20150625

Customer ID: vSRX-JuniperEval

Features:

Virtual Appliance - Virtual Appliance

count-down, Original validity: 60 days

License identifier: JUNOS657051

License version: 4

Software Serial Number: 9XXXXAXXXXXX9

Customer ID: MyCompany

Features:

Virtual Appliance - Virtual Appliance

permanent

For more information, see [Licenses for vSRX](#)