



---

## Junos OS

### Junos Node Slicing Feature Guide



---

Modified: 2018-09-18

Juniper Networks, Inc.  
1133 Innovation Way  
Sunnyvale, California 94089  
USA  
408-745-2000  
[www.juniper.net](http://www.juniper.net)

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. and/or its affiliates in the United States and other countries. All other trademarks may be property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

*Junos OS Junos Node Slicing Feature Guide*

Copyright © 2018 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://www.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 the Documentation . . . . .	xi
	Documentation and Release Notes . . . . .	xi
	Using the Examples in This Manual . . . . .	xi
	Merging a Full Example . . . . .	xii
	Merging a Snippet . . . . .	xii
	Documentation Conventions . . . . .	xiii
	Documentation Feedback . . . . .	xv
	Requesting Technical Support . . . . .	xv
	Self-Help Online Tools and Resources . . . . .	xvi
	Opening a Case with JTAC . . . . .	xvi
<b>Part 1</b>	<b>Overview</b>	
<b>Chapter 1</b>	<b>Junos Node Slicing Overview . . . . .</b>	<b>3</b>
	Junos Node Slicing Overview . . . . .	3
	Benefits of Junos Node Slicing . . . . .	4
	Components of Junos Node Slicing . . . . .	4
	Base System (BSYS) . . . . .	5
	Guest Network Function (GNF) . . . . .	5
	Juniper Device Manager (JDM) . . . . .	6
	Junos Node Slicing Administrator Roles . . . . .	6
	Mastership Behavior of BSYS and GNF . . . . .	7
	BSYS Mastership . . . . .	7
	GNF Mastership . . . . .	7
	Abstracted Fabric (AF) Interface . . . . .	8
	Understanding AF Interface Bandwidth . . . . .	8
	Features Supported on AF Interfaces . . . . .	9
	AF Interface Restrictions . . . . .	11
	Multi-Version Software Interoperability Overview . . . . .	11
	Licensing for Junos Node Slicing . . . . .	12
<b>Part 2</b>	<b>Setup</b>	
<b>Chapter 2</b>	<b>Hardware and Software Requirements for Junos Node Slicing . . . . .</b>	<b>15</b>
	Minimum Hardware and Software Requirements for Junos Node Slicing . . . . .	15
	MX Series Router . . . . .	15
	x86 Servers . . . . .	15
	Server Hardware Resource Requirements (Per GNF) . . . . .	16
	Shared Server Hardware Resource Requirements . . . . .	16
	Server Software Requirements . . . . .	17

<b>Chapter 3</b>	<b>Preparing for Junos Node Slicing Setup . . . . .</b>	<b>19</b>
	Setting Up the Connection Between Servers and the Router . . . . .	19
	x86 Server CPU BIOS Settings . . . . .	21
	x86 Server Linux GRUB Configuration . . . . .	22
	Updating Intel X710 NIC Driver for x86 Servers . . . . .	23
	Installing Additional Packages for JDM . . . . .	25
	Completing the Connection Between the Servers and the Router . . . . .	26
<b>Chapter 4</b>	<b>Setting Up Junos Node Slicing . . . . .</b>	<b>27</b>
	Completing the Prerequisite Tasks . . . . .	27
	Configuring an MX Series Router to Operate in BSYS Mode . . . . .	27
	Installing JDM RPM Package on x86 Servers Running RHEL . . . . .	29
	Installing JDM Ubuntu Package on x86 Servers Running Ubuntu 16.04 . . . . .	30
	Configuring JDM on the x86 Servers . . . . .	31
	Configuring x86 Server Interfaces in JDM . . . . .	32
	Configuring Guest Network Functions . . . . .	36
	Chassis Configuration Hierarchy at BSYS and GNF . . . . .	38
	SNMP Trap Support: Configuring NMS Server . . . . .	39
	Configuring Abstracted Fabric Interfaces Between a Pair of GNFs . . . . .	41
	Class of Service on Abstracted Fabric Interfaces . . . . .	42
	Forwarding Class-to-Queue Mapping . . . . .	42
	BA Classifiers and Rewrites . . . . .	42
	Sample Configuration for Junos Node Slicing . . . . .	43
	Sample JDM Configuration . . . . .	43
	Sample BSYS Configuration with Abstracted Fabric (AF) Interface . . . . .	45
	Sample AF Configuration at GNF with Class of Service . . . . .	45
	Sample Output for Abstracted Fabric (AF) Interface State at a GNF . . . . .	47
<b>Chapter 5</b>	<b>Setting Up YANG-Based Orchestration of GNFs . . . . .</b>	<b>49</b>
	Setting Up YANG-Based Orchestration of GNFs . . . . .	49
	Installing YANG Package on BSYS . . . . .	49
	Sharing the SSH Identities Between BSYS and JDM . . . . .	50
	Configuring the BSYS to Communicate with JDM . . . . .	51
	Disabling the YANG-based Orchestration of GNFs . . . . .	51
	Supported XML RPCs: YANG-Based Abstraction of GNFs . . . . .	52
<b>Chapter 6</b>	<b>Managing Junos Node Slicing . . . . .</b>	<b>55</b>
	Junos Node Slicing Upgrade . . . . .	55
	Upgrading Junos Node Slicing . . . . .	55
	Upgrading JDM . . . . .	56
	Upgrading GNF and BSYS . . . . .	56
	Unified ISSU Support . . . . .	56
	Managing Multi-Version Software Interoperability . . . . .	57
	Viewing Software Incompatibility Alarms . . . . .	59
	Viewing Incompatibilities Between Software Versions . . . . .	59
	Restarting External Servers . . . . .	59
	Updating Host OS on the External Servers . . . . .	61
	Configuring Non-Root Users in JDM (Junos Node Slicing) . . . . .	61
	Deleting Guest Network Functions . . . . .	62
	Disabling Junos Node Slicing . . . . .	63

<b>Part 3</b>	<b>Configuration Statements and Operational Commands</b>	
<b>Chapter 7</b>	<b>Configuration Statements for BSYS</b>	<b>67</b>
	network-slices	68
	guest-network-functions	69
	gnf	70
	control-plane-bandwidth-percent (Node Slicing)	71
	description (GNF)	72
	fpcs (Node Slicing)	73
	af-name	74
	peer-gnf	75
	description (AF)	76
<b>Chapter 8</b>	<b>Operational Commands for BSYS</b>	<b>77</b>
	show chassis network-slices	78
	show chassis fpc pic-status (Node Slicing)	81
	show chassis fpc (Node Slicing)	83
	show chassis adc (Node Slicing)	85
	show chassis network-slices fpcs	87
	show system anomalies gnf-id (BSYS)	88
<b>Chapter 9</b>	<b>Operational Commands for GNF</b>	<b>91</b>
	show chassis gnf	92
	show chassis gnf	94
	show chassis hardware (GNF)	96
	show chassis fpc (GNF)	98
	show chassis fpc pic-status (GNF)	100
	show chassis adc (GNF)	101
	show interfaces (Abstracted Fabric)	102
	show system anomalies (GNF)	112
<b>Chapter 10</b>	<b>Configuration Statements for JDM</b>	<b>115</b>
	virtual-network-functions	116
	server	117
	interfaces (Junos Node Slicing)	118
	routing-options (Junos Node Slicing)	119
	system login (Junos Nose Slicing)	120
	root-login (JDM)	121
<b>Chapter 11</b>	<b>Operational Commands for JDM</b>	<b>123</b>
	Generic Guidelines for Using JDM Server Commands	123
	clear log (JDM)	125
	monitor list (JDM)	126
	monitor start (JDM)	127
	monitor stop (JDM)	129
	request server authenticate-peer-server	130
	request virtual-network-functions	131
	show virtual-network-functions	132
	show version vnf	136
	show version (JDM)	138
	show system cpu (JDM)	140

show system storage (JDM) . . . . .	142
show system memory (JDM) . . . . .	144
show system network (JDM) . . . . .	146
restart (JDM) . . . . .	150

# List of Figures

<b>Part 1</b>	<b>Overview</b>	
<b>Chapter 1</b>	<b>Junos Node Slicing Overview</b> .....	<b>3</b>
	Figure 1: GNFs on External Server .....	5
	Figure 2: Mastership Behavior of GNF and BSYS .....	7
	Figure 3: Abstracted Fabric Interface .....	8
<b>Part 2</b>	<b>Setup</b>	
<b>Chapter 3</b>	<b>Preparing for Junos Node Slicing Setup</b> .....	<b>19</b>
	Figure 4: MX2020 Router—External x86 Server Connectivity .....	20





# List of Tables

	<b>About the Documentation</b> . . . . .	<b>xi</b>
	Table 1: Notice Icons . . . . .	xiii
	Table 2: Text and Syntax Conventions . . . . .	xiii
<b>Part 1</b>	<b>Overview</b>	
<b>Chapter 1</b>	<b>Junos Node Slicing Overview</b> . . . . .	<b>3</b>
	Table 3: Supported AF-capable MPCs . . . . .	10
<b>Part 2</b>	<b>Setup</b>	
<b>Chapter 2</b>	<b>Hardware and Software Requirements for Junos Node Slicing</b> . . . . .	<b>15</b>
	Table 4: GNF Resource Template . . . . .	16
	Table 5: Shared Server Resources Requirements . . . . .	17
<b>Chapter 5</b>	<b>Setting Up YANG-Based Orchestration of GNFs</b> . . . . .	<b>49</b>
	Table 6: Supported XML RPCs to Manage GNFs . . . . .	52
<b>Chapter 6</b>	<b>Managing Junos Node Slicing</b> . . . . .	<b>55</b>
	Table 7: Predefined Login Classes . . . . .	62
<b>Part 3</b>	<b>Configuration Statements and Operational Commands</b>	
<b>Chapter 8</b>	<b>Operational Commands for BSYS</b> . . . . .	<b>77</b>
	Table 8: show chassis network-slices Output Fields . . . . .	78
	Table 9: show chassis fpc Output Fields . . . . .	83
	Table 10: show chassis adc Output Fields . . . . .	85
	Table 11: show chassis network-slices fpcs Output Fields . . . . .	87
	Table 12: show system anomalies gnf-id Output Fields . . . . .	88
<b>Chapter 9</b>	<b>Operational Commands for GNF</b> . . . . .	<b>91</b>
	Table 13: show chassis gnf Output Fields . . . . .	92
	Table 14: show chassis gnf Output Fields . . . . .	94
	Table 15: show chassis hardware Output Fields . . . . .	96
	Table 16: show chassis fpc Output Fields . . . . .	98
	Table 17: show chassis adc Output Fields . . . . .	101
	Table 18: show interfaces (Abstracted Fabric) Output Fields . . . . .	103
	Table 19: show system anomalies gnf-id Output Fields . . . . .	112
<b>Chapter 11</b>	<b>Operational Commands for JDM</b> . . . . .	<b>123</b>
	Table 20: monitor start Output Fields . . . . .	127
	Table 21: show virtual network functions Output Fields . . . . .	132
	Table 22: show system cpu Output Fields . . . . .	140

Table 23: show system storage Output Fields . . . . .	142
Table 24: show system memory Output Fields . . . . .	144
Table 25: show system network Output Fields . . . . .	146

# About the Documentation

- Documentation and Release Notes on page xi
- Using the Examples in This Manual on page xi
- Documentation Conventions on page xiii
- Documentation Feedback on page xv
- Requesting Technical Support on page xv

## Documentation and Release Notes

---

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at <https://www.juniper.net/documentation/>.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at <https://www.juniper.net/books>.

## Using the Examples in This Manual

---

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming configuration into the current candidate configuration. The example does not become active until you commit the candidate configuration.

If the example configuration contains the top level of the hierarchy (or multiple hierarchies), the example is a *full example*. In this case, use the **load merge** command.

If the example configuration does not start at the top level of the hierarchy, the example is a *snippet*. In this case, use the **load merge relative** command. These procedures are described in the following sections.

## Merging a Full Example

To merge a full example, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration example into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following configuration to a file and name the file **ex-script.conf**. Copy the **ex-script.conf** file to the **/var/tmp** directory on your routing platform.

```
system {
  scripts {
    commit {
      file ex-script.xml;
    }
  }
}
interfaces {
  fxp0 {
    disable;
    unit 0 {
      family inet {
        address 10.0.0.1/24;
      }
    }
  }
}
```

2. Merge the contents of the file into your routing platform configuration by issuing the **load merge** configuration mode command:

```
[edit]
user@host# load merge /var/tmp/ex-script.conf
load complete
```

## Merging a Snippet

To merge a snippet, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration snippet into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following snippet to a file and name the file **ex-script-snippet.conf**. Copy the **ex-script-snippet.conf** file to the **/var/tmp** directory on your routing platform.

```
commit {
  file ex-script-snippet.xml; }
```

2. Move to the hierarchy level that is relevant for this snippet by issuing the following configuration mode command:

```
[edit]
```

```
user@host# edit system scripts
[edit system scripts]
```

3. Merge the contents of the file into your routing platform configuration by issuing the **load merge relative** configuration mode command:

```
[edit system scripts]
user@host# load merge relative /var/tmp/ex-script-snippet.conf
load complete
```

For more information about the **load** command, see [CLI Explorer](#).

## Documentation Conventions

Table 1 on page xiii defines notice icons used in this guide.

Table 1: Notice Icons







Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.
	Tip	Indicates helpful information.
	Best practice	Alerts you to a recommended use or implementation.

Table 2 on page xiii defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
<b>Bold text like this</b>	Represents text that you type.	To enter configuration mode, type the <b>configure</b> command:  user@host> <b>configure</b>

Table 2: Text and Syntax Conventions (continued)

Convention	Description	Examples
Fixed-width text like this	Represents output that appears on the terminal screen.	<pre>user@host&gt; show chassis alarms</pre> <p>No alarms currently active</p>
<i>Italic text like this</i>	<ul style="list-style-type: none"> <li>Introduces or emphasizes important new terms.</li> <li>Identifies guide names.</li> <li>Identifies RFC and Internet draft titles.</li> </ul>	<ul style="list-style-type: none"> <li>A policy <i>term</i> is a named structure that defines match conditions and actions.</li> <li><i>Junos OS CLI User Guide</i></li> <li>RFC 1997, <i>BGP Communities Attribute</i></li> </ul>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	<p>Configure the machine's domain name:</p> <pre>[edit] root@# set system domain-name domain-name</pre>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"> <li>To configure a stub area, include the <b>stub</b> statement at the <b>[edit protocols ospf area area-id]</b> hierarchy level.</li> <li>The console port is labeled <b>CONSOLE</b>.</li> </ul>
< > (angle brackets)	Encloses optional keywords or variables.	<b>stub &lt;default-metric metric&gt;;</b>
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	<b>broadcast   multicast</b> <b>(string1   string2   string3)</b>
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	<b>rsvp { # Required for dynamic MPLS only</b>
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	<b>community name members [ community-ids ]</b>
Indentation and braces ( { } )	Identifies a level in the configuration hierarchy.	<pre>[edit] routing-options {   static {     route default {       nexthop address;       retain;     }   } }</pre>
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
<b>GUI Conventions</b>		
<b>Bold text like this</b>	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> <li>In the Logical Interfaces box, select <b>All Interfaces</b>.</li> <li>To cancel the configuration, click <b>Cancel</b>.</li> </ul>

Table 2: Text and Syntax Conventions (continued)

Convention	Description	Examples
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select <b>Protocols&gt;Ospf</b> .

## Documentation Feedback

We encourage you to provide feedback so that we can improve our documentation. You can use either of the following methods:

- Online feedback system—Click TechLibrary Feedback, on the lower right of any page on the [Juniper Networks TechLibrary](#) site, and do one of the following:



- Click the thumbs-up icon if the information on the page was helpful to you.
- Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.
- E-mail—Send your comments to [techpubs-comments@juniper.net](mailto:techpubs-comments@juniper.net). Include the document or topic name, URL or page number, and software version (if applicable).

## Requesting Technical Support

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

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

## Self-Help Online Tools and Resources

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

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

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

## Opening a Case with JTAC

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

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

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



## PART 1

# Overview

- [Junos Node Slicing Overview on page 3](#)



## CHAPTER 1

# Junos Node Slicing Overview

- [Junos Node Slicing Overview on page 3](#)
- [Components of Junos Node Slicing on page 4](#)
- [Junos Node Slicing Administrator Roles on page 6](#)
- [Mastership Behavior of BSYS and GNF on page 7](#)
- [Abstracted Fabric \(AF\) Interface on page 8](#)
- [Multi-Version Software Interoperability Overview on page 11](#)
- [Licensing for Junos Node Slicing on page 12](#)

## Junos Node Slicing Overview

---

Junos Node Slicing enables service providers and large enterprises to create a network infrastructure that consolidates multiple routing functions into a single physical device. It helps leverage the benefits of virtualization without compromising on performance. In particular, Junos Node Slicing enables the convergence of multiple services on a single physical infrastructure while avoiding the operational complexity involved. It provides operational, functional, and administrative separation of functions on a single physical infrastructure that enables the network to implement the same virtualization principles the compute industry has been using for years.

Using Junos Node Slicing, you can create multiple partitions in a single physical MX Series router. These partitions are referred to as guest network functions (GNFs). Each GNF behaves as an independent router, with its own dedicated control plane, data plane, and management plane. This enables you to run multiple services on a single converged MX Series router, while still maintaining operational isolation between them. You can leverage the same physical device to create parallel partitions that do not share the control plane or the forwarding plane, but only share the same chassis, space, and power.

You can also send traffic between GNFs through the switch fabric by using an Abstracted Fabric (AF) interface, a pseudo interface that behaves as a first class Ethernet interface. An AF interface facilitates routing control, data, and management traffic between GNFs.

Junos Node Slicing supports multi-version software compatibility, thereby allowing the GNFs to be independently upgraded.

## Benefits of Junos Node Slicing

- **Converged network**—With Junos Node Slicing, service providers can consolidate multiple network services, such as video edge and voice edge, into a single physical router, while still maintaining operational separation between them. You can achieve both horizontal and vertical convergence. Horizontal convergence consolidates router functions of the same layer to a single router, while vertical convergence collapses router functions of different layers into a single router.
- **Improved scalability**—Focusing on virtual routing partitions, instead of physical devices, improves the programmability and scalability of the network, enabling service providers and enterprises to respond to infrastructure requirements without having to buy additional hardware.
- **Easy risk management**—Though multiple network functions converge on a single chassis, all the functions run independently, benefiting from operational, functional, and administrative separation. Partitioning a physical system, such as Broadband Network Gateway (BNG), into multiple independent logical instances ensures that failures are isolated. The partitions do not share the control plane or the forwarding plane, but only share the same chassis, space, and power. This means failure in one partition does not cause any widespread service outage.
- **Reduced network costs**—Junos Node Slicing enables interconnection of GNFs through internal switching fabrics, which leverages Abstracted Fabric (AF) interface, a pseudo interface that represents a first class Ethernet interface behavior. With AF interface in place, companies no longer need to depend on physical interfaces to connect GNFs, resulting in significant savings.
- **Reduced time-to-market for new services and capabilities**—Each GNF can operate on a different Junos software version. This advantage enables companies to evolve each GNF at its own pace. If a new service or a feature needs to be deployed on a certain GNF, and it requires a new software release, only the GNF involved requires an update. Additionally, with the increased agility, Junos Node Slicing enables service providers and enterprises to introduce highly flexible Everything-as-a-service business model to rapidly respond to ever-changing market conditions.

### Related Documentation

- *License Enforcement*
- *Software Feature Licenses*
- *Verifying Junos OS License Installation (CLI)*
- *show system license*

---

## Components of Junos Node Slicing

Junos Node Slicing allows a single MX Series router to be partitioned to appear as multiple, independent routers. Each partition has its own Junos OS control plane, which runs as a virtual machine (VM), and a dedicated set of line cards. Each partition is called a guest network function (GNF).

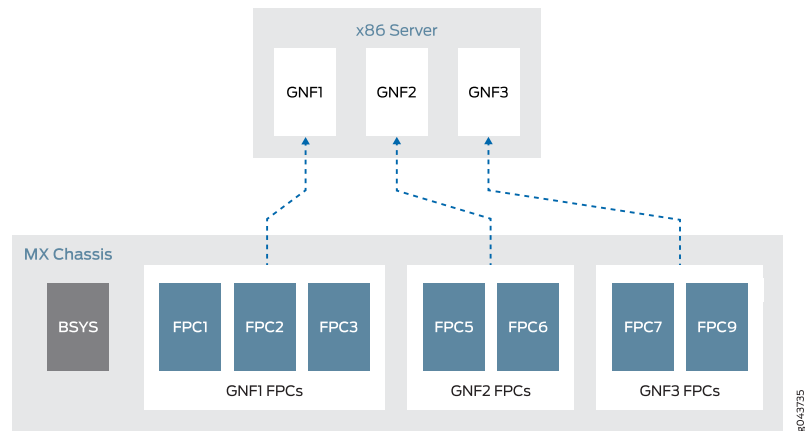
The MX Series router functions as the base system (BSYS). The BSYS owns all the physical components of the router, including the line cards and the switching fabric. The BSYS assigns line cards to GNFs.

The Juniper Device Manager (JDM) software orchestrates the GNF VMs. In JDM, a GNF VM is referred to as a virtual network function (VNF). A GNF thus comprises a VNF and a set of line cards.

JDM and VNFs are hosted on a pair of external industry standard x86 servers.

Through configuration at the BSYS, you can assign line cards of the chassis to different GNFs. [Figure 1 on page 5](#) shows three GNFs with their dedicated line cards running on an external server.

*Figure 1: GNFs on External Server*



See [“Setting Up the Connection Between Servers and the Router” on page 19](#) for information about how to connect an MX Series router to a pair of external x86 servers.

## Base System (BSYS)

In Junos Node Slicing, the MX Series router functions as the base system (BSYS). The BSYS owns all the physical components of the router, including all line cards and fabric. Through Junos OS configuration at the BSYS, you can assign line cards to GNFs and define Abstracted Fabric (AF) interfaces between GNFs. The BSYS software runs on a pair of redundant Routing Engines of the MX Series router.

## Guest Network Function (GNF)

A guest network function (GNF) logically owns the line cards assigned to it by the base system (BSYS), and maintains the forwarding state of the line cards. You can configure multiple GNFs on an MX Series router (see [“Configuring Guest Network Functions” on page 36](#)). The Junos OS control plane of each GNF runs as a virtual machine (VM). The Juniper Device Manager (JDM) software, hosted on a pair of x86 servers, orchestrates the GNF VMs. In the JDM, the GNFs are referred to as virtual network functions (VNF).

A GNF is equivalent to a standalone router. GNFs are configured and administered independently, and are operationally isolated from each other.

Creating a GNF requires two sets of configurations, one to be performed at the BSYS, and the other at the JDM.

A GNF is defined by an ID. This ID must be the same at the BSYS and JDM.

The BSYS part of the GNF configuration comprises giving it an ID and a set of line cards.

The JDM part of the GNF configuration comprises specifying the following attributes:

- A VNF name.
- A GNF ID. This ID must be the same as the GNF ID used at the BSYS.
- The MX Series platform type.
- A Junos OS image to be used for the VNF.
- The VNF CPU and memory resource profile template.

The server resource template defines the number of dedicated CPU cores and the size of DRAM to be assigned to a GNF. For a list of predefined server resource templates available for GNFs, see the *Server Hardware Resource Requirements (Per GNF)* section in [“Minimum Hardware and Software Requirements for Junos Node Slicing” on page 15](#).

After a GNF is configured, you can access it by connecting to the virtual console port of the GNF. Using the Junos OS CLI at the GNF, you can then configure the GNF system properties such as hostname and management IP address, and subsequently access it through its management port.

## Juniper Device Manager (JDM)

The Juniper Device Manager (JDM), a virtualized Linux container, enables provisioning and management of the GNF VMs.

JDM supports Junos OS-like CLI, NETCONF for configuration and management and SNMP for monitoring.

A JDM instance is hosted on each of the x86 servers. The JDM instances are typically configured as peers that synchronize the GNF configurations: when a GNF VM is created on one server, the backup GNF VM is automatically created on the other server.

An IP address and an administrator account need to be configured on the JDM. After these are configured, you can directly log in to the JDM.

### Related Documentation

- [Junos Node Slicing Overview on page 3](#)
- [Mastership Behavior of BSYS and GNF on page 7](#)

---

## Junos Node Slicing Administrator Roles

The following administrator roles enable you to carry out the node slicing tasks:

- **BSYS administrator**—Responsible for the physical chassis, as well as for GNF provisioning (assignment of line cards to GNFs). Junos OS CLI commands are available for these tasks.
- **GNF administrator**—Responsible for configuration, operation, and management of Junos OS at the GNF. All regular Junos OS CLI commands are available to the GNF administrator for these tasks.
- **JDM administrator**—Responsible for the JDM server port configuration, and for the provisioning and life-cycle management of the GNF VMs (VNFs). JDM CLI commands are available for these tasks.

**Related Documentation**

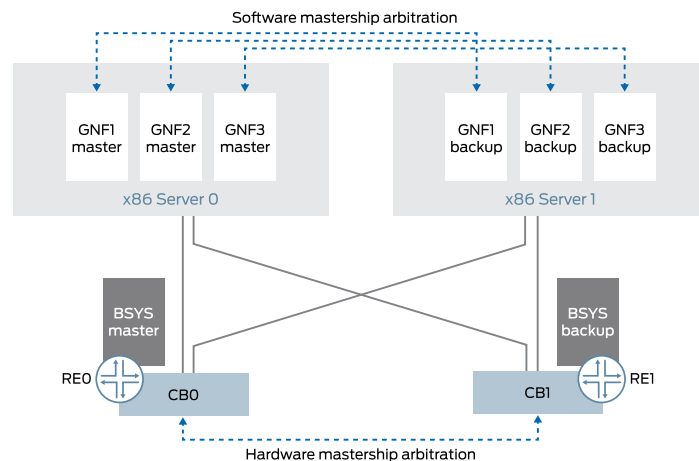
- [Components of Junos Node Slicing on page 4](#)
- [Mastership Behavior of BSYS and GNF on page 7](#)

## Mastership Behavior of BSYS and GNF

The following sections address the mastership behavior of BSYS and GNF in the context of Routing Engine redundancy.

[Figure 2 on page 7](#) shows the mastership behavior of GNF and BSYS with Routing Engine redundancy.

*Figure 2: Mastership Behavior of GNF and BSYS*



### BSYS Mastership

The BSYS Routing Engine mastership arbitration behavior is identical to that of Routing Engines on MX Series routers.

### GNF Mastership

The GNF VM mastership arbitration behavior is similar to that of MX Series Routing Engines. Each GNF runs as a master-backup pair of VMs. A GNF VM that runs on **server0**

is equivalent to Routing Engine slot 0 of an MX Series router, and the GNF VM that runs on **server1** is equivalent to Routing Engine slot 1 of an MX Series router.

The GNF mastership is independent of the BSYS mastership and that of other GNFs. The GNF mastership arbitration is done through Junos OS. Under connectivity failure conditions, GNF mastership is handled conservatively.



**NOTE:** You must configure graceful Routing Engine switchover (GRES) at each GNF. This is a prerequisite for the backup GNF VM to automatically take over the mastership when the master GNF VM fails or is rebooted.


**Related Documentation**

- [Components of Junos Node Slicing on page 4](#)

## Abstracted Fabric (AF) Interface

Abstracted fabric (AF) interface is a pseudo interface that represents a first class Ethernet interface behavior. An AF interface facilitates routing control and management traffic between guest network functions (GNFs) through the switch fabric. An AF interface is created on a GNF to communicate with its peer GNF when the two GNFs are configured to be connected to each other. AF interfaces must be created at BSYS. The bandwidth of the AF interfaces changes dynamically based on the insertion or reachability of the remote line card/MPC. Because the fabric is the communication medium between GNFs, AFs are considered to be the equivalent WAN interfaces. See [Figure 3 on page 8](#).

*Figure 3: Abstracted Fabric Interface*



## Understanding AF Interface Bandwidth

An AF interface connects two GNFs through the fabric and aggregates all the Packet Forwarding Engines (PFEs) that connect the two GNFs. An AF interface can leverage the sum of the bandwidth of each Packet Forwarding Engine belonging to the AF interface.



For example, if GNF1 has one MPC8 (which has four Packet Forwarding Engines with 240 Gbps capacity each), and GNF1 is connected with GNF2 and GNF3 using AF interfaces (af1 and af2), the maximum AF capacity of GNF1 would be  $4 \times 240 \text{ Gbps} = 960 \text{ Gbps}$ .

**GNF1—af1—GNF2**

**GNF1—af2—GNF3**

Here, af1 and af2 share the 960 Gbps capacity.

For information on the bandwidth supported on each MPC, see [Table 3 on page 10](#).

## Features Supported on AF Interfaces

AF interfaces support the following features:

- Load balancing based on the remote GNF line cards present
- Class of service (CoS) support:
  - Inet-precedence classifier and rewrite
  - DSCP classifier and rewrite
  - MPLS EXP classifier and rewrite
  - DSCP v6 classifier and rewrite for IP v6 traffic
- Support for OSPF, IS-IS, BGP, OSPFv3 protocols, and L3VPN



**NOTE:** The non-AF interfaces support all the protocols that work on Junos OS.

- Multicast forwarding
- Graceful Routing Engine switchover (GRES)
- MPLS applications where the AF interface acts as a core interface (L3VPN, VPLS, L2VPN, L2CKT, EVPN, and IP over MPLS)
- The following protocol families are supported:
  - IPv4 Forwarding
  - IPv6 Forwarding
  - MPLS
  - ISO
  - CCC
- With the AF interface configuration, GNFs support AF-capable MPCs. [Table 3 on page 10](#) lists the AF-capable MPCs, the number of PFEs supported per MPC, and the bandwidth supported per MPC.

Table 3: Supported AF-capable MPCs

MPC	Initial Release	Number of PFEs	Total Bandwidth
MPC7E-MRATE	17.4R1	2	240G (120*2) on MX240, MX480, MX960 routers 200G (100*2) on MX2010 and MX2020 routers
MPC7E-10G	17.4R1	2	240G (120*2) on MX240, MX480, MX960 routers 200G (100*2) on MX2010 and MX2020 routers
MX2K-MPC8E	17.4R1	4	960G (240*4)
MX2K-MPC9E	17.4R1	4	1.6T (400*4)
MPC2E NG	17.4R1	1	80G
MPC2E NG Q	17.4R1	1	80G
MPC3E NG	17.4R1	1	130G
MPC3E NG Q	17.4R1	1	130G
MPC5E-40G10G	18.3R1	2	240G (120*2)
MPC5EQ-40G10G	18.3R1	2	240G (120*2)
MPC5E-40G100G	18.3R1	2	240G (120*2)
MPC5EQ-40G100G	18.3R1	2	240G (120*2)
MX2K-MPC6E	18.3R1	4	520G (130*4)



## NOTE:

- A GNF that does not have the AF interface configuration supports all the MPCs that are supported by a standalone MX Series router. For the list of supported MPCs, see [MPCs Supported by MX Series Routers](#).
- We recommend that you set the MTU settings on the AF interface to align to the maximum allowed value on the XE/GE interfaces. This ensures minimal or no fragmentation of packets over the AF interface.

## AF Interface Restrictions

The following are the current restrictions of AF interfaces:

- Configurations such as single endpoint AF interface, AF interface-to-GNF mapping mismatch or multiple AF interfaces mapping to same remote GNF are not checked during commit on the BSYS. Ensure that you have the correct configurations.
- Bandwidth allocation is static, based on the MPC type.
- AF interfaces do not support the Hyper mode.
- There can be minimal traffic drops (both transit and host) during the offline/restart of an MPC hosted on a remote GNF.
- Interoperability between AF-capable MPCs and non-AF capable MPCs is not supported.

### Related Documentation

- [Configuring Abstracted Fabric Interfaces Between a Pair of GNFs on page 41](#)

## Multi-Version Software Interoperability Overview

Starting from Junos OS Release 17.4R1, Junos Node Slicing supports multi-version software compatibility, enabling the BSYS to interoperate with a guest network function (GNF) which runs a Junos OS version that is higher than the software version of the BSYS. This feature supports a range of up to two versions between GNF and BSYS. That is, the GNF software can be two versions higher than the BSYS software. Both BSYS and GNF must meet a minimum version requirement of Junos OS Release 17.4R1.



**NOTE:** The restrictions in multi-version support are also applicable to the unified ISSU upgrade process.

While JDM software versioning does not have a similar restriction with respect to the GNF or BSYS software versions, we recommend that you regularly update the JDM software. A JDM upgrade does not affect any of the running GNFs.

### Related Documentation

- [Managing Multi-Version Software Interoperability on page 57](#)
- [Junos Node Slicing Upgrade on page 55](#)
- [Restarting External Servers on page 59](#)
- [Disabling Junos Node Slicing on page 63](#)

## Licensing for Junos Node Slicing

---

Operating Junos Node Slicing requires licenses to be installed at the BSYS. Running a GNF without a license installed at the BSYS will result in the following BSYS log message:

```
CHASSISD_LICENSE_EVENT: License Network-Slices: Failed to get valid license('216')  
'gnf-creation'  
Minor alarm set, 1 Guest network functions creation for JUNOS requires a license.
```

Please contact Juniper Networks if you have queries pertaining to Junos Node Slicing licenses.

**Related Documentation**

- [Junos Node Slicing Overview on page 3](#)

## PART 2

# Setup

- [Hardware and Software Requirements for Junos Node Slicing on page 15](#)
- [Preparing for Junos Node Slicing Setup on page 19](#)
- [Setting Up Junos Node Slicing on page 27](#)
- [Setting Up YANG-Based Orchestration of GNFs on page 49](#)
- [Managing Junos Node Slicing on page 55](#)



## CHAPTER 2

# Hardware and Software Requirements for Junos Node Slicing

- [Minimum Hardware and Software Requirements for Junos Node Slicing on page 15](#)

## Minimum Hardware and Software Requirements for Junos Node Slicing

---

To set up Junos Node Slicing, you need an MX Series router and a pair of industry standard x86 servers. The x86 servers host the Juniper Device Manager (JDM) along with the GNF VMs.

### MX Series Router

The following routers support Junos Node Slicing:

- MX2010
- MX2020
- MX480
- MX960
- MX2008



**NOTE:** For the MX960 and MX480 routers, the Control Boards must be SCBE2; and the Routing Engines must be interoperable with SCBE2 (RE-S-1800, RE-S-X6-64G).

### x86 Servers

Ensure that both the servers have similar (preferably identical) hardware configuration.

The x86 server hardware resource requirements comprise:

- Per-GNF resource requirements (CPU, memory, and storage).
- Shared resource requirements (CPU, memory, storage and network ports).

The server hardware requirements are thus the sum of the requirements of the individual GNFs, and the shared resource requirements. The server hardware requirements are a function of how many GNFs you plan to use.

**x86 CPU:**

- Must be Intel Haswell-EP or newer.

**BIOS:**

- Must have hyperthreading disabled.
- Must have performance mode enabled.

**Storage:**

- Must be local to the server.
- Must be solid-state drive (SSD)-based.

The storage space for GNFs is allocated from the following:

- / (root), which must have a minimum available storage space of 50 GB.
- /vm-primary, which must have a minimum available storage space of 350 GB.

We recommend hardware RAID1 configuration for storage resiliency. We also recommend that you set up /vm-primary as a Linux partition.

### Server Hardware Resource Requirements (Per GNF)

Each GNF must be associated with a resource template, which defines the number of dedicated CPU cores and the size of DRAM to be assigned for that GNF.

Table 4 on page 16 lists the GNF resource templates available:

**Table 4: GNF Resource Template**

Template	CPU cores	DRAM (GB)
2core-16g	2	16
4core-32g	4	32
6core-48g	6	48
8core-64g	8	64

Each GNF requires a minimum of 64 GB storage.

### Shared Server Hardware Resource Requirements

Table 5 on page 17 lists the server hardware resources that are shared between all the guest network functions (GNFs) on a server:





**NOTE:** These requirements are in addition to the per-GNF requirements mentioned in the Server Hardware Resource Requirements (Per GNF) section.

*Table 5: Shared Server Resources Requirements*

Component	Specification
CPU	<ul style="list-style-type: none"> <li>Four cores to be allocated for JDM and Linux host processing.</li> </ul>
Memory	<ul style="list-style-type: none"> <li>Minimum of 32 GB DRAM for JDM and Linux host processing.</li> </ul>
Storage	<ul style="list-style-type: none"> <li>Minimum of 64 GB storage for JDM and Linux host.</li> </ul>
Network Ports	<ul style="list-style-type: none"> <li>Two 10-Gbps Ethernet interfaces for control plane connection between the server and the router.               <ul style="list-style-type: none"> <li>Minimum—1 PCIe NIC card with Intel X710 dual port 10-Gbps Direct Attach, SFP+, Converged Network Adapter, PCIe 3.0, x8</li> <li>Recommended—2 NIC cards of the above type. Use one port from each card to provide redundancy at the card level.</li> </ul> </li> <li>One Ethernet interface (1/10 Gbps) for Linux host management network.</li> <li>One Ethernet interface (1/10 Gbps) for JDM management network.</li> <li>One Ethernet interface (1/10 Gbps) for GNF management network. (This port is shared by all the GNFs on that server).</li> <li>Serial port or an equivalent interface (iDRAC, IPMI) for server console access.</li> </ul>

### Server Software Requirements

The x86 servers must have the following installed:

- Red Hat® Enterprise Linux® (RHEL) 7.3 or Ubuntu 16.04 LTS - with virtualization packages.

To enable virtualization for RHEL, choose "Virtualization Host" for the Base Environment and "Virtualization Platform" as an Add-On from the Software Selection screen during installation.



**NOTE:**

- The hypervisor supported is KVM.

- Install additional packages required for Intel X710 NIC Driver and JDM. For more information, see the ["Updating Intel X710 NIC Driver for x86 Servers" on page 23](#) and ["Installing Additional Packages for JDM" on page 25](#) sections.
- If you are using Intel X710 NIC, ensure that you have the latest driver (2.0.23 or later) installed.

The servers must also have the BIOS setup as described in ["x86 Server CPU BIOS Settings" on page 21](#) and the Linux GRUB configuration as described in ["x86 Server Linux GRUB Configuration" on page 22](#).

Ensure that the host OS is up to date.



NOTE:

- The x86 servers require internet connectivity for you to be able to perform host OS updates and install the additional packages.
- Ensure that you have the same host OS software version on both the servers.



NOTE: The following software packages are required to set up Junos Node Slicing:

- JDM package
- Junos OS image for GNFs
- Junos OS package for BSYS
- Junos OS vmhost package for REMX2K-X8-64G and RE-S-X6-64G Control Board-Routing Engine based BSYS

**Related  
Documentation**

- [Components of Junos Node Slicing on page 4](#)
- [Setting Up the Connection Between Servers and the Router on page 19](#)

## CHAPTER 3

# Preparing for Junos Node Slicing Setup

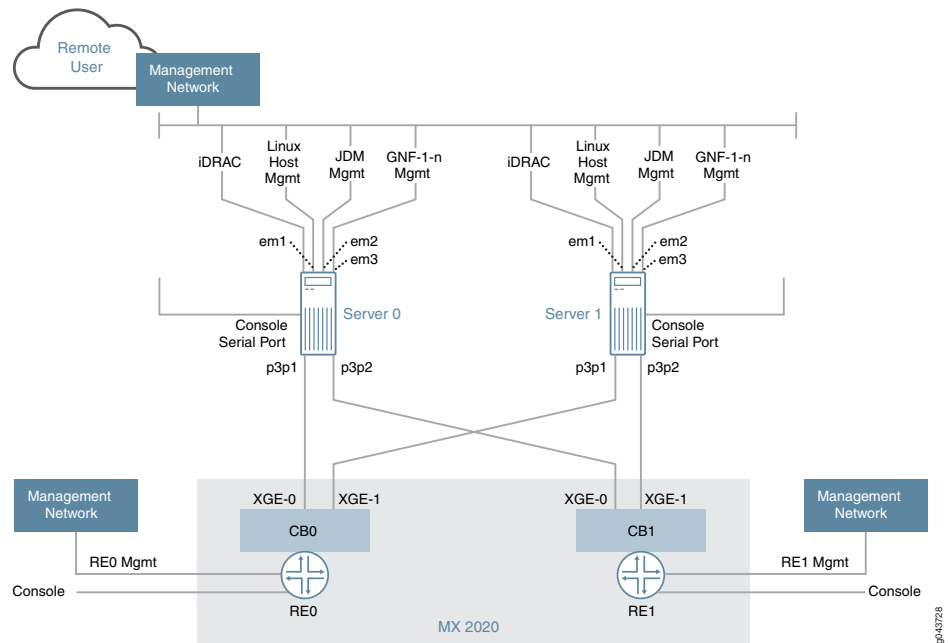
- [Setting Up the Connection Between Servers and the Router on page 19](#)
- [x86 Server CPU BIOS Settings on page 21](#)
- [x86 Server Linux GRUB Configuration on page 22](#)
- [Updating Intel X710 NIC Driver for x86 Servers on page 23](#)
- [Installing Additional Packages for JDM on page 25](#)
- [Completing the Connection Between the Servers and the Router on page 26](#)

### Setting Up the Connection Between Servers and the Router

To set up Junos Node Slicing, you must directly connect a pair of external x86 servers to the MX Series router. Besides the management port for the Linux host, each server also requires two additional ports for providing management connectivity for the JDM and the GNF VMs, respectively, and two ports for connecting to the MX Series router.

[Figure 4 on page 20](#) shows how an MX2020 router is connected to a pair of x86 external servers.

Figure 4: MX2020 Router—External x86 Server Connectivity



According to the example in [Figure 4 on page 20](#), **em1**, **em2**, and **em3** on the x86 servers are the ports that are used for the management of the Linux host, the JDM and the GNFs, respectively. **p3p1** and **p3p2** on each server are the two 10-Gbps ports that are connected to the Control Boards of the MX Series router.



**NOTE:** The names of interfaces on the server, such as **em1**, **p3p1** might vary according to the server hardware configuration.

For more information on the XGE ports of the MX Series router Control Board (CB) mentioned in [Figure 4 on page 20](#), see:

- [SCBE2-MX Description](#) (for MX960 and MX480)
- [REMX2K-X8-64G and REMX2K-X8-64G-LT CB-RE Description](#) (for MX2010 and MX2020)



**NOTE:** Use the `show chassis ethernet-switch` command to view these XGE ports. In the command output on MX960, refer to the port numbers 24 and 26 to view these ports on the SCBE2. In the command output on MX2010 and MX2020, refer to the port numbers 26 and 27 to view these ports on the Control Board-Routing Engine (CB-RE).

#### Related Documentation

- [Components of Junos Node Slicing on page 4](#)
- [Minimum Hardware and Software Requirements for Junos Node Slicing on page 15](#)

## x86 Server CPU BIOS Settings

For Junos Node Slicing, the BIOS of the x86 server CPUs should be set up such that:

- Hyperthreading is disabled.
- The CPU cores always run at their rated frequency.
- The CPU cores are set to reduce jitter by limiting C-state use.

To find the rated frequency of the CPU cores on the server, run the Linux host command **lscpu**, and check the value for the field **Model name**. See the following example:

```
Linux server0:~# lscpu
```

```
..
Model name:      Intel(R) Xeon(R) CPU E5-2680 v3 @ 2.50GHz
..
```

To find the frequency at which the CPU cores are currently running, run the Linux host command **grep MHz /proc/cpuinfo** and check the value for each CPU core.

On a server that has the BIOS set to operate the CPU cores at their rated frequency, the observed values for the CPU cores will all match the rated frequency (or be very close to it), as shown in the following example.

```
Linux server0:~# grep MHz /proc/cpuinfo
```

```
...
cpu MHz : 2499.902
cpu MHz : 2500.000
cpu MHz : 2500.000
cpu MHz : 2499.902
...
```

On a server that does not have the BIOS set to operate the CPU cores at their rated frequency, the observed values for the CPU cores do not match the rated frequency, and the values could also vary with time (you can check this by rerunning the command).

```
Linux server0:~# grep MHz /proc/cpuinfo
```

```
...
cpu MHz : 1200.562
cpu MHz : 1245.468
cpu MHz : 1217.625
cpu MHz : 1214.156
```

To set the x86 server BIOS system profile to operate the CPU cores at their rated frequency, reduce jitter, and disable hyperthreading, consult the server manufacturer, because these settings vary with server model and BIOS versions.

Typical BIOS system profile settings to achieve this include:

- **Logical processor:** set to Disabled.
- **CPU power management:** set to Maximum performance.

- **Memory frequency:** set to Maximum performance.
- **Turbo boost:** set to Disabled.
- **C-states and C1E state:** set to Disabled.
- **Energy efficient policy:** set to Performance.
- **Monitor/Mwait:** set to Disabled.

A custom BIOS system profile might be required to set these values.

**Related Documentation**

- [Minimum Hardware and Software Requirements for Junos Node Slicing on page 15](#)
- [x86 Server Linux GRUB Configuration on page 22](#)
- [Updating Intel X710 NIC Driver for x86 Servers on page 23](#)

---

## x86 Server Linux GRUB Configuration

In Junos Node Slicing, each GNF VM is assigned dedicated CPU cores. This assignment is managed by Juniper Device Manager (JDM). On each x86 server, JDM requires that all CPU cores other than CPU cores 0 and 1 be reserved for Junos Node Slicing – and in effect, that these cores be isolated from other applications. CPU cores 2 and 3 are dedicated for GNF virtual disk and network I/O. CPU cores 4 and above are available for assignment to GNF VMs. To reserve these CPU cores, you must set the **isolcpus** parameter in the Linux GRUB configuration as described in the following procedure:

For x86 servers running Red Hat Enterprise Linux (RHEL) 7.3, perform the following steps:

1. Determine the number of CPU cores on the x86 server. Ensure that hyperthreading has already been disabled, as described in [“x86 Server CPU BIOS Settings” on page 21](#). You can use the Linux command **lscpu** to find the total number of CPU cores, as shown in the following example:

```
Linux server0:~# lscpu
...
Cores per socket: 12
Sockets: 2
...
```

Here, there are 24 cores (12 x 2). The CPU cores are numbered as core 0 to core 23.

2. As per this example, the **isolcpus** parameter must be set to 'isolcpus=2-23' (isolate all CPU cores other than cores 0 and 1 for use by JDM and the GNF VMs).

To set the **isolcpus** parameter in the Linux GRUB configuration file, follow the procedure described in the section *Isolating CPUs from the process scheduler* in [this Red Hat document](#). A summary of the section is as follows:

- a. Edit the Linux GRUB file **/etc/default/grub** to append the **isolcpus** parameter to the variable GRUB\_CMDLINE\_LINUX, as shown in the following example:

```
GRUB_CMDLINE_LINUX=
"crashkernel=auto rd.lvm.lv=rhel/root rd.lvm.lv=rhel/swap rhgb quiet
isolcpus=2-23"
```

- b. Run the Linux shell command **grub2-mkconfig** to generate the updated GRUB file as shown below:

If you are using legacy BIOS, issue the following command:

```
# grub2-mkconfig -o /boot/grub2/grub.cfg
```

If you are using UEFI, issue the following command:

```
# grub2-mkconfig -o /boot/efi/EFI/redhat/grub.cfg
```

- c. Reboot the x86 server.
- d. Verify that the **isolcpus** parameter has now been set, by checking the output of the Linux command **cat /proc/cmdline**, as shown in the following example:

```
# cat /proc/cmdline
BOOT_IMAGE=/vmlinuz-3.10.0-327.36.3.el7.x86_64 ... quiet isolcpus=2-23
```

For x86 servers running Ubuntu 16.04, perform the following steps:

1. Determine the number of CPU cores on the x86 server. Ensure that hyperthreading has already been disabled, as described in x86 Server CPU BIOS Settings. You can use the Linux command **lscpu** to find the total number of CPU cores.
2. Edit the `/etc/default/grub` file to append the **isolcpus** parameter to the variable **GRUB\_CMDLINE\_LINUX\_DEFAULT**, as shown in the following example:

```
GRUB_CMDLINE_LINUX_DEFAULT=
"intel_pstate=disable processor.ignore_ppc=1 isolcpus=2-15"
```

3. To update the changes, run **update-grub**.
4. Reboot the server.
5. Verify that the **isolcpus** parameter has now been set, by checking the output of the Linux command **cat /proc/cmdline**.

#### Related Documentation

- [Minimum Hardware and Software Requirements for Junos Node Slicing on page 15](#)
- [x86 Server CPU BIOS Settings on page 21](#)

## Updating Intel X710 NIC Driver for x86 Servers

If you are using Intel X710 NIC, ensure that you have the latest driver (2.0.23 or later) installed on the x86 servers.

You need to first identify the X710 NIC interface on the servers. For example, this could be **p3p1**.

You can check the NIC driver version by running the Linux command **ethtool -i *interface***. See the following example:

```
root@Linux server0# ethtool -i p3p1

driver: i40e
version: 2.0.23
firmware-version: 5.05 0x80002899 17.5.12
...
```

Refer to the [Intel support page](#) for instructions on updating the driver.



**NOTE:** Ensure that the host OS is up to date prior to updating the Intel X710 NIC driver.

---

You need the following packages for building the driver:

- For RedHat:
  - kernel-devel
  - Development Tools
- For Ubuntu:
  - make
  - gcc

If you are using RedHat, run the following commands to install the packages:

```
root@Linux server0#yum install kernel-devel
root@Linux server0#yum group install "Development Tools"
```

If you are using Ubuntu, run the following commands to install the packages:

```
root@Linux server0# apt-get install make
root@Linux server0# apt-get install gcc
```



**NOTE:** After updating the Intel X710 NIC driver, you might notice the following message in the host OS log:

**"i40e: module verification failed: signature and/or required key missing - tainting kernel"**

Ignore this message. It appears because the updated NIC driver module has superseded the base version of the driver that was packaged with the host OS.

---



- Related Documentation**
- [Minimum Hardware and Software Requirements for Junos Node Slicing on page 15](#)

## Installing Additional Packages for JDM

The x86 servers must have Red Hat Enterprise Linux (RHEL) 7.3 or Ubuntu 16.04 LTS installed.



**NOTE:** The x86 Servers must have the virtualization packages installed.

**For RHEL 7.3**, install the following additional packages, which can be downloaded from the [Red Hat Customer Portal](#).

- python-psutil-1.2.1-1.el7.x86\_64.rpm
- net-snmp-5.7.2-24.el7.x86\_64.rpm
- net-snmp-libs-5.7.2-24.el7.x86\_64.rpm
- libvirt-snmp-0.0.3-5.el7.x86\_64.rpm

Only for Junos OS Releases 17.4R1 and earlier, and for 18.1R1, if you are running RHEL 7.3, also install the following additional package:

- libstdc++-4.8.5-11.el7.i686.rpm



**NOTE:**

- The package version numbers shown are the minimum versions. Newer versions might be available in the latest RHEL 7.3 patches.
- The libstdc++ package extension .i686 indicates that it is a 32-bit package.
- For RHEL, we recommend that you install the packages using the yum command.

**For Ubuntu 16.04**, install the following packages:

- python-psutil

Only for Junos OS Releases 17.4R1 and earlier, and for 18.1R1, if you are running Ubuntu, also install the following additional package:

- `libstdc++6:i386`

**NOTE:**

- For Ubuntu, you can use the `apt-get` command to install the latest version of these packages. For example, use:
  - the command `apt-get install python-psutil` to install the latest version of the `python-psutil` package.
  - the command `apt-get install libstdc++6:i386` to install the latest version of the `libstdc++6` package (the extension `:i386` indicates that the package being installed is a 32-bit version).

**Related Documentation**

- [Components of Junos Node Slicing on page 4](#)
- [Minimum Hardware and Software Requirements for Junos Node Slicing on page 15](#)
- [Setting Up the Connection Between Servers and the Router on page 19](#)
- [Sample Configuration for Junos Node Slicing on page 43](#)

---

## Completing the Connection Between the Servers and the Router

---

Complete the following steps before you start installing the JDM:

- Ensure that the MX Series router is connected to the x86 servers as described in “[Setting Up the Connection Between Servers and the Router](#)” on page 19.
- Power on the two x86 servers and both the Routing Engines on the MX Series router.
- Identify the Linux host management port on both the x86 servers. For example, `em1`.
- Identify the ports to be assigned for the JDM and the GNF management ports. For example, `em2` and `em3`.
- Identify the two 10-Gbps ports that are connected to the Control Boards on the MX Series router. For example, `p3p1` and `p3p2`.

**Related Documentation**

- [Components of Junos Node Slicing on page 4](#)
- [Minimum Hardware and Software Requirements for Junos Node Slicing on page 15](#)
- [Setting Up the Connection Between Servers and the Router on page 19](#)
- [Sample Configuration for Junos Node Slicing on page 43](#)

## CHAPTER 4

# Setting Up Junos Node Slicing

- [Completing the Prerequisite Tasks on page 27](#)
- [Configuring an MX Series Router to Operate in BSYS Mode on page 27](#)
- [Installing JDM RPM Package on x86 Servers Running RHEL on page 29](#)
- [Installing JDM Ubuntu Package on x86 Servers Running Ubuntu 16.04 on page 30](#)
- [Configuring JDM on the x86 Servers on page 31](#)
- [Configuring x86 Server Interfaces in JDM on page 32](#)
- [Configuring Guest Network Functions on page 36](#)
- [Chassis Configuration Hierarchy at BSYS and GNF on page 38](#)
- [SNMP Trap Support: Configuring NMS Server on page 39](#)
- [Configuring Abstracted Fabric Interfaces Between a Pair of GNFs on page 41](#)
- [Sample Configuration for Junos Node Slicing on page 43](#)

### Completing the Prerequisite Tasks

---

Before proceeding to perform the Junos Node Slicing setup tasks, you must have completed the procedures described in the chapter 'Preparing for Junos Node Slicing Setup'.

#### Related Documentation

- [Minimum Hardware and Software Requirements for Junos Node Slicing on page 15](#)

### Configuring an MX Series Router to Operate in BSYS Mode

---



NOTE: Ensure that the MX Series router is connected to the x86 servers as described in “[Setting Up the Connection Between Servers and the Router](#)” on page 19.

Junos Node Slicing requires the MX Series router to function as the base system (BSYS).

Use the following steps to configure an MX Series router to operate in BSYS mode:

1. Install the Junos OS package for BSYS on both the Routing Engines of the MX Series router.

To download the package:

- a. Go to the [Juniper support page for Junos Node Slicing](#).
  - b. Click **Base System** > *Junos OS version number* > *Junos version number (64-bit High-End)*.
  - c. On the **Software Download** page, select the **I Agree** option under End User License Agreement and then click **Proceed**.
2. On the MX Series router, run the **show chassis hardware** command and verify that the transceivers on both the Control Boards (CBs) are detected. The following text represents a sample output:

```
root@router> show chassis hardware
```

```
...
CB 0          REV 23   750-040257   CABL4989   Control Board
  Xcvr 0      REV 01   740-031980   ANT00F9   SFP+-10G-SR
  Xcvr 1      REV 01   740-031980   APG0SC3   SFP+-10G-SR
CB 1          REV 24   750-040257   CABX8889   Control Board
  Xcvr 0      REV 01   740-031980   AP41BKS   SFP+-10G-SR
  Xcvr 1      REV 01   740-031980   ALN0PCM   SFP+-10G-SR
```

3. On the MX Series router, apply the following configuration statements:

```
root@router# set chassis network-slices guest-network-functions
root@router# set chassis redundancy graceful-switchover
root@router# set chassis network-services enhanced-ip
root@router# set routing-options nonstop-routing
root@router# set system commit synchronize
root@router# commit
```



**NOTE:** On MX960 routers, you must configure the network-services mode as enhanced-ip or enhanced-ethernet. On MX2020 routers, the enhanced-ip configuration statement is already enabled by default.

The router now operates in BSYS mode.



**NOTE:** A router in the BSYS mode is not expected to run features other than the ones required to run the basic management functionalities in Junos Node Slicing. For example, the BSYS is not expected to have interface configurations associated with the line cards installed in the system. Instead, guest network functions (GNFs) will have the full-fledged router configurations.

**Related Documentation**

- [Components of Junos Node Slicing on page 4](#)
- [Minimum Hardware and Software Requirements for Junos Node Slicing on page 15](#)
- [Setting Up the Connection Between Servers and the Router on page 19](#)
- [Sample Configuration for Junos Node Slicing on page 43](#)

## Installing JDM RPM Package on x86 Servers Running RHEL

Before installing the JDM RPM package for x86 servers, ensure that you have installed the additional packages, as described in ["Installing Additional Packages for JDM" on page 25](#).

Download and install the JDM RPM package for x86 servers running RHEL as follows:

To download the package:

- Go to the [Juniper support page for Junos Node Slicing](#).
- Click **JDM > Junos OS version number > Juniper Device Manager version number (for Redhat)**.
- On the **Software Download** page, select the **I Agree** option under the End User License Agreement and then click **Proceed**.

To install the package on x86 servers running RHEL, perform the following steps on each of the servers:

- Disable SELINUX and reboot the server. You can disable SELINUX by setting the value for **SELINUX** to **disabled** in the `/etc/selinux/config` file.
- Install the JDM RPM package (indicated by the `.rpm` extension) by using the following command. An example of the JDM RPM package used is shown below:

```
root@Linux Server0# rpm -ivh jns-jdm-1.0-0-17.4R1.73.x86_64.rpm
```

```
Preparing...                               ##### [100%]
Detailed log of jdm setup saved in /var/log/jns-jdm-setup.log
Updating / installing...
 1:jns-jdm-1.0-0                               ##### [100%]
Setup host for jdm...
Launch libvirtd in listening mode
Done Setup host for jdm
Installing /juniper/.tmp-jdm-install/juniper_ubuntu_rootfs.tgz...
Configure /juniper/lxc/jdm/jdm1/rootfs...
Configure /juniper/lxc/jdm/jdm1/rootfs DONE
Created symlink from /etc/systemd/system/multi-user.target.wants/jdm.service
to /usr/lib/systemd/system/jdm.service.
Done Setup jdm
Redirecting to /bin/systemctl restart rsyslog.service
```

Repeat the steps for the second server.

**Related Documentation**

- [Components of Junos Node Slicing on page 4](#)
- [Minimum Hardware and Software Requirements for Junos Node Slicing on page 15](#)
- [Setting Up the Connection Between Servers and the Router on page 19](#)
- [Sample Configuration for Junos Node Slicing on page 43](#)

---

## Installing JDM Ubuntu Package on x86 Servers Running Ubuntu 16.04

---

Before installing the JDM Ubuntu package for x86 servers, ensure that you have installed the additional packages. For more details, see "[Installing Additional Packages for JDM on page 25](#)."

Download and install the JDM Ubuntu package for x86 servers running Ubuntu 16.04 as follows:

To download the JDM Ubuntu package:

- a. Go to the [Juniper support page for Junos Node Slicing](#).
- b. Click **JDM** > *Junos OS version number* > *Juniper Device Manager version number (for Debian)*.
- c. On the **Software Download** page, select the **I Agree** option under the End User License Agreement and then click **Proceed**.

To install the JDM package on the x86 servers running Ubuntu 16.04, perform the following steps on each of the servers:

1. Disable **apparmor** and reboot the server.

```
root@Linux Server0# systemctl stop apparmor
```

```
root@Linux Server0# systemctl disable apparmor
```

```
root@Linux Server0# reboot
```

2. Install the JDM Ubuntu package (indicated by the *.deb* extension) by using the following command. An example of the JDM Ubuntu package used is shown below:

```
root@Linux Server0# dpkg -i jns-jdm-1.0-0-17.4R1.13.x86_64.deb
Selecting previously unselected package jns-jdm.
(Reading database ... 71846 files and directories currently installed.)
Preparing to unpack jns-jdm-1.0-0-17.4R1.13.x86_64.deb ...
Unpacking jns-jdm (1.0-0) ...
Setting up jns-jdm (1.0-0) ...
Installing /juniper/.tmp-jdm-install/juniper_ubuntu_latest.tgz...
Configure /juniper/lxc/jdm/jdm1/rootfs...
Configure /juniper/lxc/jdm/jdm1/rootfs DONE
Done Setup jdm
Processing triggers for ureadahead (0.100.0-19) ...
Processing triggers for systemd (225-1ubuntu9) ...
```

Repeat the steps for the second server.

**Related Documentation**

- [Components of Junos Node Slicing on page 4](#)
- [Minimum Hardware and Software Requirements for Junos Node Slicing on page 15](#)
- [Setting Up the Connection Between Servers and the Router on page 19](#)
- [Sample Configuration for Junos Node Slicing on page 43](#)

## Configuring JDM on the x86 Servers

Use the following steps to configure JDM on each of the x86 servers.

1. At each server, start the JDM, and assign identities for the two servers as **server0** and **server1**, respectively, as follows:

On one server, run the following command:

```
root@Linux server0# jdm start server=0
```

Starting JDM

On the other server, run the following command:

```
root@Linux server1# jdm start server=1
```

Starting JDM



**NOTE:** The identities, once assigned, cannot be modified without uninstalling the JDM and then reinstalling it:

2. Enter the JDM console on each server by running the following command:

```
root@Linux Server0# jdm console
```

```
Connected to domain jdm
```

```
Escape character is ^]
```

```
* Starting Signal sysvinit that the rootfs is mounted [ OK ]
* Starting Populate /dev filesystem [ OK ]
* Starting Populate /var filesystem [ OK ]
* Stopping Send an event to indicate plymouth is up [ OK ]
* Stopping Populate /var filesystem [ OK ]
* Starting Clean /tmp directory [ OK ]
```

```
...
```

```
jdm login:
```

3. Log in as the **root** user.

4. Enter the JDM CLI by running the following command:

```
root@jdm% cli
```



**NOTE:** The JDM CLI is similar to the Junos OS CLI.

5. Set the root password for the JDM.

```
root@jdm# set system root-authentication plain-text-password
```

New Password:



**NOTE:**

- The JDM supports root user administration account only.
- The JDM root password must be the same on both the servers.
- The JDM root password overrides the Linux host root password.

6. Commit the changes:

```
root@jdm# commit
```

7. Enter **Ctrl-]** to exit from the JDM console.

8. From the Linux host, run the **ssh jdm** command to log in to the JDM shell.

**Related  
Documentation**

- [Components of Junos Node Slicing on page 4](#)
- [Minimum Hardware and Software Requirements for Junos Node Slicing on page 15](#)
- [Setting Up the Connection Between Servers and the Router on page 19](#)
- [Sample Configuration for Junos Node Slicing on page 43](#)

---

## Configuring x86 Server Interfaces in JDM

In the JDM, you must configure:

- The two 10-Gbps server ports that are connected to the MX Series router.
- The server port to be used as the JDM management port.
- The server port to be used as the GNF management port.

Therefore, you need to identify the following on each server before starting the configuration of the ports:

- The server interfaces (for example, **p3p1** and **p3p2**) that are connected to **CB0** and **CB1** on the MX Series router.



- The server interfaces (for example, **em2** and **em3**) to be used for JDM management and GNF management.

For more information, see the figure [Figure 4 on page 20](#).



**NOTE:**

- You need this information for both **server0** and **server1**.
- These interfaces are visible only on the Linux host.

To configure the x86 server interfaces in JDM, perform the following steps on both the servers:

1. On **server0**, apply the following configuration statements:

```
root@jdm# set groups server0 server interfaces cb0 p3p1
root@jdm# set groups server0 server interfaces cb1 p3p2
root@jdm# set groups server1 server interfaces cb0 p3p1
root@jdm# set groups server1 server interfaces cb1 p3p2
root@jdm# set apply-groups [ server0 server1 ]
root@jdm# commit

root@jdm# set groups server0 server interfaces jdm-management em2
root@jdm# set groups server0 server interfaces vnf-management em3
root@jdm# set groups server1 server interfaces jdm-management em2
root@jdm# set groups server1 server interfaces vnf-management em3
root@jdm# commit
```

2. Repeat the step 1 on **server1**.



**NOTE:** Ensure that you apply the same configuration on both **server0** and **server1**.

3. Share the **ssh** identities between the two x86 servers.

At both **server0** and **server1**, run the following JDM CLI command:

```
root@jdm> request server authenticate-peer-server
```



**NOTE:** The request server authenticate-peer-server command displays a CLI message requesting you to log in to the peer server using ssh to verify the operation. To log in to the peer server, you need to prefix ip netns exec jdm\_nv\_ns to ssh root@jdm-server1.

For example, to log in to the peer server from server0, exit the JDM CLI, and use the following command from JDM shell:

```
root@jdm:~# ip netns exec jdm_nv_ns ssh root@jdm-server1
```

Similarly, to log in to the peer server from server1, use the following command:

```
root@jdm:~# ip netns exec jdm_nv_ns ssh root@jdm-server0
```

4. Apply the configuration statements in the JDM CLI configuration mode to set the JDM management IP address, default route, and the JDM hostname for each JDM instance as shown in the following example.



**NOTE:** The management IP address and default route must be specific to your network.

```
root@jdm# set groups server0 interfaces jmgmt0 unit 0 family inet address 10.216.105.112/21
root@jdm# set groups server1 interfaces jmgmt0 unit 0 family inet address 10.216.105.113/21
root@jdm# set groups server0 routing-options static route 0.0.0.0/0 next-hop 10.216.111.254
root@jdm# set groups server1 routing-options static route 0.0.0.0/0 next-hop 10.216.111.254
root@jdm# set groups server0 system host-name test-jdm-server0
root@jdm# set groups server1 system host-name test-jdm-server1
root@jdm# commit
```



**NOTE:**

- jmgmt0 stands for the JDM management port. This is different from the Linux host management port. Both JDM and the Linux host management ports are independently accessible from the management network.

5. Run the following JDM CLI command on each server and ensure that all the interfaces are up.

```
root@jdm> show server connections
```

Status	Comments	Component	Interface
up		Host to JDM port	virbr0
up		Physical CB0 port	p3p1
up		Physical CB1 port	p3p2

up	Physical JDM mgmt port	em2
up	Physical VNF mgmt port	em3
up	JDM-GNF bridge	bridge_jdm_vm
up	CB0	cb0
up	CB1	cb1
up	JDM mgmt port	jmgmt0
up	JDM to HOST port	bme1
up	JDM to GNF port	bme2
up	JDM to JDM link0*	cb0.4002
up	JDM to JDM link1	cb1.4002



**NOTE:** For sample JDM configurations, see [“Sample Configuration for Junos Node Slicing”](#) on page 43.

If you want to modify the server interfaces configured in the JDM, perform the following steps:

1. Stop all running GNFs.

```
root@jdm> request virtual-network-functions gnf-name stop
```

2. From the configuration mode, deactivate the virtual network functions configuration, and then commit the change.

```
root@jdm# deactivate virtual-network-functions
root@jdm# commit
```

3. Configure and commit the new interfaces as described in the step 1 of the main procedure.

4. Reboot the JDM from the shell.

```
root@jdm:~# reboot
```

5. From the configuration mode, activate the virtual network functions configuration, and then commit the change.

```
root@jdm# activate virtual-network-functions
root@jdm# commit
```

**Related Documentation**

- [Components of Junos Node Slicing on page 4](#)
- [Minimum Hardware and Software Requirements for Junos Node Slicing on page 15](#)
- [Setting Up the Connection Between Servers and the Router on page 19](#)
- [Sample Configuration for Junos Node Slicing on page 43](#)

---

## Configuring Guest Network Functions

---

Configuring a guest network function (GNF) comprises two tasks, one to be performed at the BSYS and the other at the JDM.



---

**NOTE:**

- Before attempting to create a GNF, you must ensure that the servers have sufficient resources (CPU, memory, storage) for that GNF.
  - You need to assign an ID to each GNF. This ID must be the same at the BSYS and the JDM.
- 

At the BSYS, specify a GNF by assigning it an ID and a set of line cards by applying the configuration as shown in the following example:

```
user@router# set chassis network-slices guest-network-functions gnf 1 fpcs 4  
  
user@router# commit
```

In the JDM, the GNF VMs are referred to as virtual network functions (VNFs). A VNF has the following attributes:

- A VNF name.
- A GNF ID. This ID must be the same as the GNF ID used at the BSYS.
- The MX Series platform type.
- A Junos OS image to be used for the GNF.
- The VNF CPU and memory resource profile template..

To configure a VNF, perform the following steps:

1. Retrieve the *Junos OS image for GNFs* and place it in the host OS directory `/var/jdm-usr/gnf-images/` on both the servers.

To download the package:

- a. Go to the [Juniper support page for Junos Node Slicing](#).
  - b. Click **GNF** > *Junos OS version number* > *Junos version number (Guest Network Function)* .
  - c. On the **Software Download** page, select the **I Agree** option under the End User License Agreement and then click **Proceed**.
2. Assign this image to a GNF by using the JDM CLI command as shown in the following example:

```
root@test-jdm-server0> request virtual-network-functions test-gnf add-image
/var/jdm-usr/gnf-images/junos-install-ns-mx-x86-64-17.4R1.10.tgz all-servers
Server0:
Added image: /vm-primary/test-gnf/test-gnf.img
Server1:
Added image: /vm-primary/test-gnf/test-gnf.img
```

3. Configure the VNF by applying the configuration statements as shown in the following example:

```
root@test-jdm-server0# set virtual-network-functions test-gnf id 1

root@test-jdm-server0# set virtual-network-functions test-gnf chassis-type mx2020

root@test-jdm-server0# set virtual-network-functions test-gnf resource-template
2core-16g
```

To also specify a baseline or initial Junos OS configuration for a GNF, prepare the GNF configuration file (example: `/var/jdm-usr/gnf-config/test-gnf.conf`) on both the servers and specify the filename as the parameter in the **base-config** statement as shown below:

```
root@test-jdm-server0# set virtual-network-functions test-gnf base-config
/var/jdm-usr/gnf-config/test-gnf.conf

root@test-jdm-server0# commit synchronize
```



NOTE: Ensure that:

- You use the same GNF ID as the one specified earlier in BSYS.
- The baseline configuration filename (with the path) is the same on both the servers.
- The syntax of the baseline file contents is in the Junos OS configuration format.
- The GNF name used here is the same as the one assigned to the Junos OS image for GNF in the step 2.

4. To verify that the VNF is created, run the following JDM CLI command:

```
root@test-jdm-server0> show virtual-network-functions test-gnf
```

5. Log in to the console of the VNF by issuing the following JDM CLI command:

```
root@test-jdm-server0> request virtual-network-functions test-gnf console
```

6. Configure the VNF the same way as you configure an MX Series Routing Engine.



NOTE:

- For sample configurations, see [“Sample Configuration for Junos Node Slicing” on page 43](#).
- If you had previously brought down any physical x86 CB interfaces or the GNF management interface from Linux shell (by using the command `ifconfig interface-name down`), these will automatically be brought up when the GNF is started.

#### Related Documentation

- [Components of Junos Node Slicing on page 4](#)
- [Minimum Hardware and Software Requirements for Junos Node Slicing on page 15](#)
- [Setting Up the Connection Between Servers and the Router on page 19](#)
- [Sample Configuration for Junos Node Slicing on page 43](#)

---

## Chassis Configuration Hierarchy at BSYS and GNF

In Junos Node Slicing, the BSYS owns all the physical components of the router, including the line cards and fabric, while the GNFs maintain forwarding state on their respective line cards. In keeping with this split responsibility, Junos CLI configuration under the **chassis** hierarchy (if any), should be applied at the BSYS or at the GNF as follows:

- Physical-level parameters under the **chassis** configuration hierarchy should be applied at the BSYS. For example, the configuration for handling physical errors at an FPC is a physical-level parameter, and should therefore be applied at the BSYS.

At BSYS Junos CLI:

[edit]

```
user@router# set chassis fpc fpc slot error major threshold threshold value action alarm
```

- Logical or feature-level parameters under the **chassis** configuration hierarchy should be applied at the GNF associated with the FPC. For example, the configuration for max-queues per line card is a logical-level parameter, and should therefore be applied at the GNF.

At GNF Junos CLI:

[edit]

```
user@router# set chassis fpc fpc slot max-queues value
```

- As exceptions, the following two parameters under the **chassis** configuration hierarchy should be applied at both BSYS and GNF:

At both BSYS and GNF CLI:

[edit]

```
user@router# set chassis network-services network services mode
```

```
user@router# set chassis fpc fpc slot flexible-queueing-mode
```

Related Documentation • [Configuring Guest Network Functions on page 36](#)

## SNMP Trap Support: Configuring NMS Server

The Juniper Device Manager (JDM) supports the following SNMP traps:

- LinkUp and linkDown traps for JDM interfaces.

Standard linkUp/linkDown SNMP traps are generated. A default community string **jdm** is used.

- LinkUp/linkDown traps for host interfaces.

Standard **linkUp/linkDown** SNMP traps are generated. A default community string **host** is used.

- JDM to JDM connectivity loss/regain traps.

JDM to JDM connectivity loss/regain traps are sent using generic syslog traps (jnxSyslogTrap) through the host management interface.

The JDM connectivity down trap **JDM\_JDM\_LINK\_DOWN** is sent when the JDM is not able to communicate with the peer JDM on another server over **cb0** or **cb1** links. See the following example:

```
{ SNMPv2c C=host { V2Trap(296) R=1299287309
.1.3.6.1.2.1.1.3.0=42761992
.1.3.6.1.6.3.1.1.4.1.0=.1.3.6.1.4.1.2636.4.12.0.1
```

```
.1.3.6.1.4.1.2636.3.35.1.1.1.2.1="JDM_JDM_LINK_DOWN"
.1.3.6.1.4.1.2636.3.35.1.1.1.3.1=""
.1.3.6.1.4.1.2636.3.35.1.1.1.4.1=5
.1.3.6.1.4.1.2636.3.35.1.1.1.5.1=24
.1.3.6.1.4.1.2636.3.35.1.1.1.6.1=0
.1.3.6.1.4.1.2636.3.35.1.1.1.7.1="jdmmon"
.1.3.6.1.4.1.2636.3.35.1.1.1.8.1="JDM-HOST"
.1.3.6.1.4.1.2636.3.35.1.1.1.9.1="JDM to JDM Connection Lost"
.1.3.6.1.6.3.1.1.4.3.0.0="" } }
```

The JDM to JDM Connectivity up trap **JDM\_JDM\_LINK\_UP** is sent when either the **cb0** or **cb1** link comes up, and JDMs on both the servers are able to communicate again. See the following example:

```
{ SNMPv2c C=host { V2Trap(292) R=998879760
.1.3.6.1.2.1.1.3.0=42762230
.1.3.6.1.6.3.1.1.4.1.0=.1.3.6.1.4.1.2636.4.12.0.1
.1.3.6.1.4.1.2636.3.35.1.1.1.2.1="JDM_JDM_LINK_UP"
.1.3.6.1.4.1.2636.3.35.1.1.1.3.1=""
.1.3.6.1.4.1.2636.3.35.1.1.1.4.1=5
.1.3.6.1.4.1.2636.3.35.1.1.1.5.1=24
.1.3.6.1.4.1.2636.3.35.1.1.1.6.1=0
.1.3.6.1.4.1.2636.3.35.1.1.1.7.1="jdmmon"
.1.3.6.1.4.1.2636.3.35.1.1.1.8.1="JDM-HOST"
.1.3.6.1.4.1.2636.3.35.1.1.1.9.1="JDM to JDM Connection Up"
.1.3.6.1.6.3.1.1.4.3.0.0="" } }
```

- VM(GNF) up/down—**libvirtGuestNotif** notifications.

For GNF start/shutdown events, the standard **libvirtGuestNotif** notifications are generated. For **libvirtMIB** notification details, see this [web page](#). Also, see the following example:

```
HOST [UDP: [127.0.0.1]:53568->[127.0.0.1]]: Trap ,
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (636682) 1:46:06.82,
SNMPv2-MIB::snmpTrapOID.0 = OID: LIBVIRT-MIB::libvirtGuestNotif,
LIBVIRT-MIB::libvirtGuestName.0 = STRING: "gnf1",
LIBVIRT-MIB::libvirtGuestUUID.1 = STRING: 7ad4bc2a-16db-d8c0-1f5a-6cb777e17cd8,
LIBVIRT-MIB::libvirtGuestState.2 = INTEGER: running(1),
LIBVIRT-MIB::libvirtGuestRowStatus.3 = INTEGER: active(1)
```

SNMP traps are sent to the target NMS server. To configure the target NMS server details in the JDM, see the following example:

[edit]

```
root@jdm# show snmp | display set
root@jdm# set snmp name name
root@jdm# set snmp description description
root@jdm# set snmp location location
root@jdm# set snmp contact user's email
root@jdm# set snmp trap-group tg-1 targets target ip address1
root@jdm# set snmp trap-group tg-1 targets target ip address2
```

## Related Documentation

- [Installing JDM RPM Package on x86 Servers Running RHEL on page 29](#)



- [Installing JDM Ubuntu Package on x86 Servers Running Ubuntu 16.04 on page 30](#)
- [Configuring JDM on the x86 Servers on page 31](#)

## Configuring Abstracted Fabric Interfaces Between a Pair of GNFs

Creating an Abstracted Fabric (AF) interface between two guest network functions (GNFs) involves configurations both at the base system (BSYS) and at the GNF. AF interfaces are created on GNFs based on the BSYS configuration, which is then sent to those GNFs.



**NOTE:** Only one AF interface can be configured between a pair of GNFs.

To configure AF interfaces between a pair of GNFs:

1. At the BSYS, apply the configuration as shown in the following example:

```
user@router# set chassis network-slices guest-network-functions gnf 2 af4 peer-gnf id 4
user@router# set chassis network-slices guest-network-functions gnf 2 af4 peer-gnf af2
user@router# set chassis network-slices guest-network-functions gnf 4 af2 peer-gnf id 2
user@router# set chassis network-slices guest-network-functions gnf 4 af2 peer-gnf af4
```

In this example, **af2** is the Abstracted Fabric interface instance 2 and **af4** is the Abstracted Fabric interface instance 4.



**NOTE:** The allowed AF interface values range from **af0** through **af9**.

The GNF AF interface will be visible and up. You can configure an AF interface the way you configure any other interface.

2. At the GNF, apply the configuration as shown in the following example:

```
user@router-gnf-b# set interfaces af4 unit 0 family inet address 10.10.10.1/24
user@router-gnf-d# set interfaces af2 unit 0 family inet address 10.10.10.2/24
```



**NOTE:**

- If you want to apply MPLS family configurations on the AF interfaces, you can apply the command `set interfaces af-name unit logical-unit-number family mpls` on both the GNFs between which the AF interface is configured.
- For sample AF configurations, see [“Sample Configuration for Junos Node Slicing” on page 43](#).

## Class of Service on Abstracted Fabric Interfaces

Class of service (CoS) packet classification assigns an incoming packet to an output queue based on the packet's forwarding class. See [CoS Configuration Guide](#) for more details.

The following sections explain the forwarding class- to-queue mapping, and the behavior aggregate (BA) classifiers and rewrites supported on the Abstracted Fabric (AF) interfaces.

---

### Forwarding Class-to-Queue Mapping

An AF interface is a simulated WAN interface with most capabilities of any other interface except that the traffic designated to a remote Packet Forwarding Engine will still have to go over the two fabric queues (Low/High priority ones).



**NOTE:** Presently, the AF interface operates in 2-queue mode only. Hence, all queue-based features such as scheduling, policing, and shaping are not available on an AF interface.

---

Packets on the AF interface inherit the fabric queue that is determined by the fabric priority configured for the forwarding class to which that packet belongs. For example, see the following forwarding class to queue map configuration:

[edit]

```
user@router# show class-of-service forwarding-classes
```

```
class Economy queue-num 0 priority low; /* Low fabric priority */
class Stream queue-num 1;
class Business queue-num 2;
class Voice queue-num 3;
class NetControl queue-num 3;
class Business2 queue-num 4;
class Business3 queue-num 5;
class VoiceSig queue-num 6 priority high; /* High fabric priority */
class VoiceRTP queue-num 7;
```

As shown in the preceding example, when a packet gets classified to the forwarding class **VoiceSig**, the code in the forwarding path examines the fabric priority of that forwarding class and decides which fabric queue to choose for this packet. In this case, high-priority fabric queue is chosen.

---

### BA Classifiers and Rewrites

The behavior aggregate (BA) classifier maps a class-of-service (CoS) value to a forwarding class and loss priority. The forwarding class and loss-priority combination determines the CoS treatment given to the packet in the router. The following BA classifiers and rewrites are supported:

- Inet-Precedence classifier and rewrite
- DSCP classifier and rewrite
- MPLS EXP classifier and rewrite

You can also apply rewrites for IP packets entering the MPLS tunnel and do a rewrite of both EXP and IPv4 type of service (ToS) bits. This approach will work as it does on other normal interfaces.

- DSCP v6 classifier and rewrite for IP v6 traffic



#### NOTE:

The following are not supported:

- IEEE 802.1 classification and rewrite
- IEEE 802.1AD (QinQ) classification and rewrite

See [CoS Configuration Guide](#) for details on CoS BA classifiers.

#### Related Documentation

- [Abstracted Fabric \(AF\) Interface on page 8](#)

## Sample Configuration for Junos Node Slicing

This section provides sample configurations for Junos Node Slicing.

### Sample JDM Configuration

```
root@test-jdm-server0> show configuration
groups {
server0 {
  system {
    host-name test-jdm-server0;
  }
  server {
    interfaces {
      cb0 p3p1;
      cb1 p3p2;
      jdm-management em2;
      vnf-management em3;
    }
  }
}
interfaces {
  jmgmt0 {
    unit 0 {
      family inet {
        address 10.216.105.112/21;
      }
    }
  }
}
}
```

```
routing-options {
  static {
    route {
      0.0.0.0/0 next-hop 10.216.111.254;
    }
  }
}
server1 {
  system {
    host-name test-jdm-server1;
  }
  server {
    interfaces {
      cb0 p3p1;
      cb1 p3p2;
      jdm-management em2;
      vnf-management em3;
    }
  }
  interfaces {
    jmgmt0 {
      unit 0 {
        family inet {
          address 10.216.105.113/21;
        }
      }
    }
  }
  routing-options {
    static {
      route {
        0.0.0.0/0 next-hop 10.216.111.254;
      }
    }
  }
}
}
apply-groups [ server0 server1 ];
system {
  root-authentication {
    encrypted-password "..."; ## SECRET-DATA
  }
  services {
    ssh;
    netconf {
      ssh;
      rfc-compliant;
    }
  }
}
virtual-network-functions {
  test-gnf {
    id 1;
    chassis-type mx2020;
    resource-template 2core-16g;
```

```

        base-config /var/jdm-usr/gnf-config/test-gnf.conf;
    }
}

```

## Sample BSYS Configuration with Abstracted Fabric (AF) Interface

```

user@router> show configuration chassis
network-slices {
  guest-network-functions {
    gnf 1 {
      af2 {
        peer-gnf id 2 af1;
      }
      af4 {
        peer-gnf id 4 af1;
      }
      description gnf-a;
      fpcs [ 0 19];
    }
    gnf 2 {
      af1 {
        peer-gnf id 1 af2;
      }
      af4 {
        peer-gnf id 4 af2;
      }
      description gnf-b;
      fpcs [ 1 6 ];
    }
    gnf 4 {
      af1 {
        peer-gnf id 1 af4;
      }
      af2 {
        peer-gnf id 2 af4;
      }
      description gnf-d;
      fpcs [ 3 4 ];
    }
  }
}

```

## Sample AF Configuration at GNF with Class of Service

Assume that there is an AF interface between GNF1 and GNF2. The following sample configuration illustrates how to apply rewrites on the AF interface at GNF1 and apply classifiers on the AF interface on GNF2, in a scenario where traffic comes from GNF1 to GNF2:

### GNF1 Configuration

```

interfaces {
  xe-4/0/0 {
    unit 0 {
      family inet {
        address 22.1.2.2/24;
      }
    }
  }
}

```

```
    }
  }
}
af2 {
  unit 0 {
    family inet {
      address 32.1.2.1/24;
    }
  }
}
}
class-of-service {
  classifiers {
    dscp testdscp {
      forwarding-class assured-forwarding {
        loss-priority low code-points [ 001001 000000 ];
      }
    }
  }
}
interfaces {
  xe-4/0/0 {
    unit 0 {
      classifiers {
        dscp testdscp;
      }
    }
    classifiers {
      dscp testdscp;
    }
  }
}
af1 {
  unit 0 {
    rewrite-rules {
      dscp testdscp; /*Rewrite rule applied on egress AF interface on GNF1.*/
    }
  }
}
}
rewrite-rules {
  dscp testdscp {
    forwarding-class assured-forwarding {
      loss-priority low code-point 001001;
    }
  }
}
}
```

#### GNF2 Configuration

```
interfaces {
  xe-3/0/0:0 {
    unit 0 {
      family inet {
        address 42.1.2.1/24;
      }
    }
  }
}
```

```

}
af1 {
  unit 0 {
    family inet {
      address 32.1.2.2/24;
    }
  }
}
}
class-of-service {
  classifiers {
    dscp testdscp {
      forwarding-class network-control {
        loss-priority low code-points 001001;
      }
    }
  }
}
interfaces {
  af1 {
    unit 0 {
      classifiers {
        dscp testdscp; /*Classifier applied on AF at ingress of GNF2*/
      }
    }
  }
}
}

```

### Sample Output for Abstracted Fabric (AF) Interface State at a GNF

user@router-gnf-b> show interfaces af1

Physical interface: af1, Enabled, Physical link is Up  
 Interface index: 137, SNMP ifIndex: 520  
 Description: To-Other-GNF  
 Type: Ethernet, Link-level type: Ethernet, MTU: 9192, Speed: 3200000mbps  
 Device flags : Present Running  
 Interface flags: Internal: 0x4000  
 Link type : Full-Duplex  
 Link flags : None  
 Current address: 00:05:86:1b:6a:31, Hardware address: 00:05:86:1b:6a:31  
 Last flapped : 2017-06-21 20:24:59 EDT (20:20:23 ago)  
 Input rate : 648 bps (0 pps)  
 Output rate : 3225773096 bps (816237 pps)  
 Bandwidth : 3200 Gbps  
 Peer GNF id : 1

Peer GNF Forwarding element(FE) view :

FPC slot:FE Num	FE Bandwidth(Gbps)	Status
0:0	400	Up
0:1	400	Up
0:2	400	Up
0:3	400	Up
19:0	400	Up
19:1	400	Up
19:2	400	Up
19:3	400	Up

Logical interface af1.0 (Index 353) (SNMP ifIndex 663)

```

Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
Input packets : 118336
Output packets: 25759170901
Protocol inet, MTU: 9100
Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1, Curr new hold
cnt: 0, NH drop cnt: 0
  Flags: Sendbroadcast-pkt-to-re, User-MTU
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 145.2.1.14/31, Local: 145.2.1.15
  Protocol iso, MTU: 9175
  Flags: User-MTU
  Protocol inet6, MTU: 9178
  Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1, Curr new hold
cnt: 0, NH drop cnt: 0
    Addresses, Flags: Is-Preferred
      Destination: fe80::/64, Local: fe80::21f:6c0f:fc9b:8fdd
    Protocol mpls, MTU: 9174, Maximum labels: 5
    Flags: User-MTU
    Protocol multiservice, MTU: Unlimited

```



## CHAPTER 5

# Setting Up YANG-Based Orchestration of GNFs

- [Setting Up YANG-Based Orchestration of GNFs on page 49](#)
- [Disabling the YANG-based Orchestration of GNFs on page 51](#)
- [Supported XML RPCs: YANG-Based Abstraction of GNFs on page 52](#)

## Setting Up YANG-Based Orchestration of GNFs

---

Starting from Junos OS Release 17.4R1, Junos supports YANG-based abstraction to orchestrate guest network functions (GNFs), using single touchpoint.

In the single touchpoint method, a software-defined networking (SDN) controller (for example, OpenDaylight or ODL), communicates with the base system (BSYS)—that is, the BSYS acts as the single touchpoint. The BSYS receives the RPC requests from the ODL controller, parses the RPC and then forwards the adequate RPC to the JDM (based on scripts available at the BSYS). On getting the response from the JDM, the BSYS parses and forwards the response back to ODL.



**NOTE:** Junos Node Slicing also supports GNF life cycle management using the dual touchpoint method. In this method, ODL sends RPCs to, and receive responses from, JDM and BSYS separately. To enable dual touchpoint, you need to mount both BSYS and Juniper Device Manager (JDM) on ODL.

To enable single touchpoint, you need to perform the following steps.

- [Installing YANG Package on BSYS on page 49](#)
- [Sharing the SSH Identities Between BSYS and JDM on page 50](#)
- [Configuring the BSYS to Communicate with JDM on page 51](#)

## Installing YANG Package on BSYS

In the node slicing setup, two servers (referred to as Server0 and Server1) are connected to an MX series router (or the BSYS). On the BSYS, you need to install the YANG package to enable communication between the BSYS and the JDM.

To install the YANG package on the BSYS, use the following command:

```
root@router> request system software add /var/tmp/junos-node-slicing-x86-32-17.4.tgz
```



**NOTE:** You need to install the same package on the backup Routing Engine as well.

After successful installation, you can find the YANG Package contents at the following location:

```
root@router:/opt/yang-pkg/junos-node-slicing
```

You can verify the package installation status on both the Routing Engines of the BSYS, using the following command:

```
root@router> show system yang package
```

## Sharing the SSH Identities Between BSYS and JDM

To synchronize configurations between the BSYS and the JDM (on server0 and server1), the SSH identities need to be shared between the BSYS and the JDM. It is mandatory to share the authorization keys for passwordless authentication from BSYS to JDM.



**NOTE:**

- Ensure that server0 and server1 are up and running.
- Ensure that you replace the jdm-server0-ip and jdm-server1-ip with proper IP addresses.

Do the following operation on BSYS for JDM on server0 and server1:

1. Go to the BSYS shell.

```
root@router> start shell
```

2. Issue the following commands:

```
root@router:~ # ssh-keygen -t rsa
root@router:~ # ssh root@jdm-server0-ip mkdir -p .ssh
root@router:~ # cat .ssh/id_rsa.pub | ssh root@jdm-server0-ip 'cat >> .ssh/authorized_keys'
root@router:~ # ssh root@jdm-server0-ip "chmod 700 .ssh; chmod 640 .ssh/authorized_keys"
```

```
root@router:~ # ssh-keygen -t rsa
root@router:~ # ssh root@jdm-server1-ip mkdir -p .ssh
root@router:~ # cat .ssh/id_rsa.pub | ssh root@jdm-server1-ip 'cat >> .ssh/authorized_keys'
root@router:~ # ssh root@jdm-server1-ip "chmod 700 .ssh; chmod 640 .ssh/authorized_keys"
```

## Configuring the BSYS to Communicate with JDM



**NOTE:** This is a one-time setup.

To complete the single touchpoint setup, you need to set the BSYS to communicate with the JDM. This enables the BSYS to service the RPC requests from ODL.

To enable the BSYS to communicate with the JDM, issue the following commands on the BSYS.

```
root@router# set jdm-configuration:jdm jdm-management server0 jdm-server0-ip-addr
root@router# set jdm-configuration:jdm jdm-management server1 jdm-server1-ip-addr
root@router# set system scripts commit file jdm-config-check.slax
root@router# set event-options policy JNS events UI_COMMIT_COMPLETED
root@router# set event-options policy JNS then event-script jdm-config-commit.slax
root@router# set event-options event-script file jdm-config-commit.slax
root@router# commit and-quit
```

### Related Documentation

- [Disabling the YANG-based Orchestration of GNFs on page 51](#)
- [Supported XML RPCs: YANG-Based Abstraction of GNFs on page 52](#)

## Disabling the YANG-based Orchestration of GNFs

If you want to disable the YANG-based orchestration of GNFs, delete the YANG package. Also, in case of a software upgrade, you need to delete the existing package and install the latest one.

To delete the YANG package:

1. Issue the following commands on the BSYS:

```
root@router# delete jdm-configuration-jdm
root@router# delete system scripts commit file jdm-config-check.slax
root@router# delete event-options policy JNS events UI_COMMIT_COMPLETED
root@router# delete event-options policy JNS then event-script jdm-config-commit.slax
root@router# delete event-options event-script file jdm-config-commit.slax
root@router# commit and-quit
```

2. Issue the following command on both the master and backup Routing Engines:

```
root@router> request system software delete junos-node-slicing
```

### Related Documentation

- [Setting Up YANG-Based Orchestration of GNFs on page 49](#)

## Supported XML RPCs: YANG-Based Abstraction of GNFs

Table 6 on page 52 lists the key GNF management tasks, along with the XML RPCs that are used to perform those tasks.



### NOTE:

- In the single touchpoint method, use the XML with the prefix `jdm-`, as seen in Table 6 on page 52 (for example, `<rpc><jdm-get-route-information/></rpc>`).
- In the dual touchpoint method, use the XML without the `jdm-` prefix (for example, `<rpc><get-route-information/></rpc>`).

Table 6: Supported XML RPCs to Manage GNFs

Task	RPC
Add Junos OS package for a GNF.	<pre> &lt;rpc&gt;   &lt;jdm-request-virtual-network-functions&gt;     &lt;vnf-name&gt;       GNF-NAME     &lt;/vnf-name&gt;     &lt;add-image&gt;       IMAGE-PATH     &lt;/add-image&gt;     &lt;all-servers/&gt;   &lt;/jdm-request-virtual-network-functions&gt; &lt;/rpc&gt; </pre>
Delete Junos OS package of a GNF.	<pre> &lt;rpc&gt;   &lt;jdm-request-virtual-network-functions&gt;     &lt;vnf-name&gt;       GNF-NAME     &lt;/vnf-name&gt;     &lt;delete-image/&gt;     &lt;server0/&gt;   &lt;/jdm-request-virtual-network-functions&gt; &lt;/rpc&gt; </pre>
Show information about a specific GNF.	<pre> &lt;rpc&gt;   &lt;jdm-get-virtual-network-functions&gt;     &lt;vnf-name&gt;       GNF-NAME     &lt;/vnf-name&gt;     &lt;server1/&gt;   &lt;/jdm-get-virtual-network-functions&gt; &lt;/rpc&gt; </pre>

Table 6: Supported XML RPCs to Manage GNFs (continued)

Task	RPC
Stop a GNF.	<pre> &lt;rpc&gt;   &lt;jdm-request-virtual-network-functions&gt;     &lt;vnf-name&gt;       GNF-NAME     &lt;/vnf-name&gt;     &lt;stop/&gt;   &lt;/jdm-request-virtual-network-functions&gt; &lt;/rpc&gt; </pre>
Start a GNF.	<pre> &lt;rpc&gt;   &lt;jdm-request-virtual-network-functions&gt;     &lt;vnf-name&gt;       GNF-NAME     &lt;/vnf-name&gt;     &lt;start/&gt;   &lt;/jdm-request-virtual-network-functions&gt; &lt;/rpc&gt; </pre>
Restart a GNF.	<pre> &lt;rpc&gt;   &lt;jdm-request-virtual-network-functions&gt;     &lt;vnf-name&gt;       GNF-NAME     &lt;/vnf-name&gt;     &lt;restart/&gt;   &lt;/jdm-request-virtual-network-functions&gt; &lt;/rpc&gt; </pre>
Get information about the JDM software.	<pre> &lt;rpc&gt;    &lt;jdm-get-software-information/&gt; &lt;/rpc&gt; </pre>
Get route information.	<pre> &lt;rpc&gt;   &lt;jdm-get-route-information/&gt; &lt;/rpc&gt; </pre>
Get information about JDM server connections.	<pre> &lt;rpc&gt;    &lt;jdm-get-server-connections/&gt; &lt;/rpc&gt; </pre>
Get the list of JDM bridges.	<pre> &lt;rpc&gt;   &lt;jdm-get-bridges-list/&gt; &lt;/rpc&gt; </pre>
Get information about JDM interfaces.	<pre> &lt;rpc&gt;   &lt;jdm-get-interface-information/&gt; &lt;/rpc&gt; </pre>

**Table 6: Supported XML RPCs to Manage GNFs (continued)**

Get list of GNFs.

```
<rpc>  
  
<jdm-get-inventory-software-vnf-information/>  
</rpc>
```

Get the GNF visibility details.

```
<rpc>  
  
<jdm-get-visibility-vnf-information/>  
</rpc>
```

---

**Related Documentation** • [Setting Up YANG-Based Orchestration of GNFs on page 49](#)

## CHAPTER 6

# Managing Junos Node Slicing

- [Junos Node Slicing Upgrade on page 55](#)
- [Managing Multi-Version Software Interoperability on page 57](#)
- [Restarting External Servers on page 59](#)
- [Updating Host OS on the External Servers on page 61](#)
- [Configuring Non-Root Users in JDM \(Junos Node Slicing\) on page 61](#)
- [Deleting Guest Network Functions on page 62](#)
- [Disabling Junos Node Slicing on page 63](#)

### Junos Node Slicing Upgrade

---

- [Upgrading Junos Node Slicing on page 55](#)
- [Unified ISSU Support on page 56](#)

### Upgrading Junos Node Slicing

Junos Node Slicing comprises three types of software components:

- Juniper Device Manager (JDM) package
- Junos OS image for guest network function (GNFs)
- Junos OS package for base system (BSYS)

You can upgrade each of these components independently, as long as they are within the allowed range of software versions (see [“Multi-Version Software Interoperability Overview” on page 11](#) for more details). You can also upgrade all of them together.



**NOTE:** Before starting the upgrade process, save the JDM, GNF VM, and BSYS configurations for reference.

- [Upgrading JDM on page 56](#)
- [Upgrading GNF and BSYS on page 56](#)

## Upgrading JDM

---

1. Upgrade the JDM by performing the following tasks on both the servers:
  - a. Copy the new JDM package (RPM or Ubuntu) to a directory on the host (for example, `/var/tmp`).
  - b. Stop the JDM by using the following command:

```
root@Linux server0# jdm stop
Stopping JDM
```

- c. Issue the upgrade command to upgrade the JDM package:

If you are upgrading the JDM RHEL package, use the following command:

```
root@Linux server0# rpm -U package_name.rpm --force
```

If you are upgrading the JDM Ubuntu package, use the following command:

```
root@Linux server0# dpkg -i deb package.deb
```



**NOTE:** A JDM upgrade does not affect any of the running GNFs.

---

See also:

- [Installing JDM RPM Package on x86 Servers Running RHEL on page 29](#)
- [Installing JDM Ubuntu Package on x86 Servers Running Ubuntu 16.04 on page 30](#)

## Upgrading GNF and BSYS

---

The GNF and BSYS packages can be upgraded in the same way as you would upgrade Junos OS on a standalone MX Series router.



**NOTE:** You can also overwrite an existing GNF image with a new one through JDM by using the JDM CLI command `request virtual-network-functions vnf-name add-image new-image-name force`. This can be useful in a rare situation where the GNF image does not boot. You can also use the `force` option to perform a cleanup if, for example, you abruptly stopped an earlier `add-image` process by pressing Ctrl-C (example: `request virtual-network-functions vnf-name delete-image image-name force`).

---

## Unified ISSU Support

Junos Node Slicing also supports unified in-service software upgrade (ISSU), enabling you to upgrade between two different Junos OS versions with no disruption on the control plane and with minimal disruption of traffic. You can perform unified ISSU on BSYS and



GNFs separately. Also, you can run unified ISSU on each GNF independently—without affecting other GNFs. See also [Understanding the Unified ISSU Process](#).



**NOTE:** The multi-version software support restrictions (such as version deviation limits) are applicable to unified ISSU upgrade as well.

#### Related Documentation

- [Multi-Version Software Interoperability Overview on page 11](#)
- [Restarting External Servers on page 59](#)
- [Deleting Guest Network Functions on page 62](#)
- [Disabling Junos Node Slicing on page 63](#)

## Managing Multi-Version Software Interoperability

Junos Node Slicing supports multi-version software interoperability. However, if there are any incompatibilities between software versions, alert messages appear during the software upgrade process or when a GNF or a FRU comes online. When minor incompatibilities occur, you can choose to accept them and proceed. In case of a major incompatibility, you need to either abort the process or use the **force** option to accept the incompatibility and proceed.

The following are sample messages that appear if incompatibilities are detected during software upgrade:

#### Sample alert message indicating a minor incompatibility:

```
user@router> request system software add
/var/tmp/junos-install-mx-x86-64-17.4-20170703_dev_common.0.tgz
Starting Multiversion compatibility checks for package
/var/tmp/junos-install-mx-x86-64-17.4-20170703_dev_common.0.tgz
Starting compatibility checks...
```

-----  
System Anomalies:

-----  
Ano-ID   ACTION   MESSAGE  
-----

100    WARN   <sample system incompatibility 1>

Accept incompatibility? [yes,no] (no) yes

103    WARN   <sample system incompatibility 2>

Accept incompatibility? [yes,no] (no) yes

-----  
CFG Anomalies for: set snmp interface

-----  
FRU-ID   Ano-ID   ACTION   MESSAGE  
-----

NONE        102    WARN   <sample config incompatibility 1>

```
Accept incompatibility? [yes,no] (no) yes

NONE          105    WARN  <sample config incompatibility 2>

Accept incompatibility? [yes,no] (no) yes
```

```
-----
FRU Anomalies:
-----
```

```
FRU-ID  Ano-ID  ACTION  MESSAGE
-----
```

```
0xaa0b      100    WARN  <sample FRU incompatibility 1>
```

```
Accept incompatibility? [yes,no] (no) yes
```

```
0xbb0b      101    WARN  <sample FRU incompatibility 2>
```

```
Accept incompatibility? [yes,no] (no) yes
```

```
Compatibility Checks done... OK
NOTICE: Validating configuration against
junos-install-mx-x86-64-17.4-20170703_dev_common.0.tgz.
NOTICE: Use the 'no-validate' option to skip this if desired.
Verified junos-install-mx-x86-64-17.4-20170703_dev_common.0 signed by
PackageDevelopmentEc_2017 method ECDSA256+SHA256
```

#### Sample alert message indicating a major incompatibility:

```
user@router> request system software add
/var/tmp/junos-install-mx-x86-64-17.4I20170713_0718.tgz
Starting Multiversion compatibility checks for package
/var/tmp/junos-install-mx-x86-64-17.4I20170713_0718.tgz
Starting compatibility checks...
-----
System Anomalies:
-----
Ano-ID  ACTION  MESSAGE
-----
1677721600  ABORT  <sample system incompatibility 1>

error: Junos-Node-Slicing multi-version checks returned abort for package
/var/tmp/junos-install-mx-x86-64-17.4I20170713_0718.tgz
```

#### Sample output showing how to use the 'force' option to proceed with an upgrade:

```
user@router> request system software add
/var/tmp/junos-install-mx-x86-64-17.4I20170713_0718.tgz force

NOTICE: Validating configuration against
junos-install-mx-x86-64-17.4I20170713_0718.tgz.
NOTICE: Use the 'no-validate' option to skip this if desired.
Verified junos-install-mx-x86-64-17.4I20170713_0718 signed by
PackageDevelopmentEc_2017 method ECDSA256+SHA256
Verified manifest signed by PackageDevelopmentEc_2017 method ECDSA256+SHA256
Checking PIC combinations
Adding junos-x86-64-17.4I20170713_0718...
```

- [Viewing Software Incompatibility Alarms on page 59](#)
- [Viewing Incompatibilities Between Software Versions on page 59](#)

## Viewing Software Incompatibility Alarms

After a software update of a GNF or BSYS, if software incompatibilities between the GNF and the BSYS exist, they will be raised as a chassis alarm. You can view the incompatibility alarm information by using the **show chassis alarms** command. You can further view the details of the incompatibilities by using the **show system anomalies** command. For more details, see [“Viewing Incompatibilities Between Software Versions” on page 59](#).

The alarms appear only on GNFs even if the upgrade is performed on the BSYS. The following types of alarm can occur:

- **System Incompatibility with BSYS**—This is a major alarm. It appears when any incompatibilities between BSYS and GNF software versions cause the GNF to go offline.
- **Feature Incompatibility with BSYS**—This is a minor alarm. It indicates a minor incompatibility between BSYS and GNF software versions. This does not cause the GNF to go offline.

## Viewing Incompatibilities Between Software Versions

To view software incompatibilities from the BSYS, use the CLI as shown in the following example:

```
user@router> show system anomalies gnf-id 4 system
```

To view software incompatibilities from a GNF, use the CLI as shown in the following example:

```
user@router> show system anomalies system
```



### NOTE:

- As shown in the CLI, remember to specify the GNF ID while viewing the incompatibilities from BSYS.
- The preceding examples show system-level incompatibilities. Use the **fru** or **config** options to view FRU or feature-level incompatibilities.

### Related Documentation

- [Junos Node Slicing Upgrade on page 55](#)
- [Multi-Version Software Interoperability Overview on page 11](#)

## Restarting External Servers

Server maintenance activities such as hardware or host OS upgrade and fault isolation might require you to restart the external servers used in Junos Node Slicing. Use the following procedure to restart the servers:

1. Stop all the GNFs.

If you are restarting both the servers, choose the **all-servers** option while stopping each GNF as shown in the following example:

```
root@server1> request virtual-network-functions gnf_name stop all-servers
gnf_name stopped
```

If you are restarting a particular server, stop the GNFs on that server by specifying the server-id as shown in the following example:

```
root@server1> request virtual-network-functions gnf_name stop server0
gnf_name stopped
```

2. Verify that the GNFs have been stopped.

```
root@server1> show virtual-network-functions
```

ID	Name	State	Liveness
1	mgb-gnf-b	Shutdown	down



**NOTE:** If you want to view the status of GNFs on both the servers, choose the **all-servers** option. Example: `show virtual-network-functions all-servers`).

3. From the Linux host shell, stop the JDM by using the following command:

```
[root@HostLinux ~]# jdm stop
Stopping JDM
```

4. From the Linux host shell, verify that the JDM status shows as stopped.

```
[root@HostLinux ~]# jdm status
JDM is stopped
```

5. After rebooting, verify that the JDM status now shows as running.

```
[root@HostLinux ~]# jdm status
JDM (pid 2828) is running as server1
```

After a server reboot, the JDM and the configured GNFs will automatically start running.

If you are replacing the servers, ensure that the operating server pair continues to have similar or identical hardware configuration. If the server pair were to become temporarily dissimilar during the replacement (this could be the case when replacing the servers sequentially), it is recommended that you disable GRES and NSR for this period, and re-enable them only when both the servers are similar once again.

#### Related Documentation

- [Junos Node Slicing Upgrade on page 55](#)

## Updating Host OS on the External Servers

Before updating the host OS on an external server, you must first stop the GNFs and JDM on that server as described in [“Restarting External Servers” on page 59](#).

Following the host OS update, if you are using Intel X710 NICs, ensure that the version of the X710 NIC driver in use continues to be the latest version as described in [“Updating Intel X710 NIC Driver for x86 Servers” on page 23](#).

- Related Documentation**
- [Junos Node Slicing Overview on page 3](#)
  - [Junos Node Slicing Upgrade on page 55](#)

## Configuring Non-Root Users in JDM (Junos Node Slicing)

You can create non-root users on Juniper Device Manager (JDM) for Junos Node Slicing, starting in Junos OS Release 18.3R1. You need a root account to create a non-root user. The non-root users can log in to JDM by using the JDM console or through SSH. Each non-root user is provided a username and assigned a predefined login class.

The non-root users can perform the following functions:

- Interact with JDM.
- Orchestrate and manage Guest Network Functions (GNFs).
- Monitor the state of the JDM, the host server and the GNFs by using JDM CLI commands.



**NOTE:** The non-root user accounts function only inside JDM, not on the host server.

To create non-root users in JDM:

1. Log in to JDM as a root user.

2. Define a user name and assign the user with a predefined login class.

```
root@jdm# set system login user username class predefined-login-class
```

3. Set the password for the user.

```
root@jdm# set system login user username authentication plain-text-password
```

New Password:

4. Commit the changes.

```
root@jdm# commit
```

Table 7 on page 62 contains the predefined login classes that JDM supports for non-root users:

Table 7: Predefined Login Classes

Login Class	Permissions
super-user	<ul style="list-style-type: none"> <li>• Create, delete, start and stop GNFs.</li> <li>• Start and stop daemons inside the JDM.</li> <li>• Execute all CLIs.</li> <li>• Access the shell.</li> </ul>
operator	<ul style="list-style-type: none"> <li>• Start and stop GNFs.</li> <li>• Restart daemons inside the JDM.</li> <li>• Execute all basic CLI operational commands (except the ones which modify the GNFs or JDM configuration).</li> </ul>
read-only	Similar to operator class, except that the users cannot restart daemons inside JDM.
unauthorized	Ping and traceroute operations

- Related Documentation**
- [system login \(Junos Node Slicing\) on page 120](#)
  - [Configuring JDM on the x86 Servers on page 31](#)

## Deleting Guest Network Functions

This procedure involves shutting down a GNF and then deleting it. In JDM, GNF VMs are called VNFs. Use the following steps to delete a VNF:

1. Shut down a VNF by using the JDM CLI command **request virtual-network-functions *gnf-name* stop all-servers**. For example:
 

```
root@test-jdm-server0> request virtual-network-functions test-gnf stop all-servers

server0:
-----
test-gnf stopped

server1:
-----
test-gnf stopped
```
2. Delete the VNF configuration by applying the JDM CLI configuration statement **delete virtual-network-functions *gnf-name***. See the following example:
 

```
root@test-jdm-server0# delete virtual-network-functions test-gnf
root@test-jdm-server0# commit synchronize
```
3. Delete the VNF image repository by using the JDM CLI command **request virtual-network-functions *gnf-name* delete-image all-servers**. For example:

```
root@test-jdm-server0> request virtual-network-functions test-gnf delete-image
all-servers
```

```
server0:
```

```
-----
Deleted the image repository
```

```
server1:
```

```
-----
Deleted the image repository
.....
```

**NOTE:**

- To delete a VNF completely, you must perform all the three steps.
- If you want to delete a VNF management interface, you must stop and delete the VNF first.

**Related Documentation**

- [Components of Junos Node Slicing on page 4](#)
- [Sample Configuration for Junos Node Slicing on page 43](#)

## Disabling Junos Node Slicing

To disable Junos Node Slicing, you must uninstall the following packages:

- JDM package
- Junos OS image for GNFs



**NOTE:** Save the JDM configuration if you want to use it for reference.

Use the following steps to disable Junos Node Slicing:

1. Delete the GNFs first by performing all the steps described in the section [“Deleting Guest Network Functions” on page 62](#).

2. Stop the JDM on each server by running the following command at the host Linux shell:

```
root@Linux server0# jdm stop
```

```
Stopping jdm: Domain jdm destroyed
```

3. Uninstall the JDM on each server by running the following command at the host Linux shell.

For the servers running RHEL, run the following command:

```
root@Linux server0# rpm -e jns-jdm
```

For the servers running Ubuntu, run the following command:

```
root@Linux server0# dpkg --remove jns-jdm
```

4. To revert the MX Series router from BSYS mode to standalone mode, apply the following configuration statements on the MX Series router:

```
root@router# delete chassis network-slices guest-network-functions
root@router# commit
```

The router now operates in standalone mode.

**Related  
Documentation**

- [Components of Junos Node Slicing on page 4](#)
- [Sample Configuration for Junos Node Slicing on page 43](#)



## PART 3

# Configuration Statements and Operational Commands

- [Configuration Statements for BSYS on page 67](#)
- [Operational Commands for BSYS on page 77](#)
- [Operational Commands for GNF on page 91](#)
- [Configuration Statements for JDM on page 115](#)
- [Operational Commands for JDM on page 123](#)



## CHAPTER 7

# Configuration Statements for BSYS

- [network-slices on page 68](#)
- [guest-network-functions on page 69](#)
- [gnf on page 70](#)
- [control-plane-bandwidth-percent \(Node Slicing\) on page 71](#)
- [description \(GNF\) on page 72](#)
- [fpcs \(Node Slicing\) on page 73](#)
- [af-name on page 74](#)
- [peer-gnf on page 75](#)
- [description \(AF\) on page 76](#)

## network-slices

---

**Syntax**    network-slices {  
              guest-network-functions {  
                  gnf *id* {  
                      control-plane-bandwidth-percent *percent*;  
                      description *description*;  
                      fpcs *fpcs*;  
                      af-name  
                  }  
              }  
          }

**Hierarchy Level**    [edit chassis]

**Release Information**    Statement introduced in Junos OS Release 17.2.

**Description**    Configure Junos Node Slicing.

Junos Node Slicing enables a single MX Series router to be partitioned to appear as multiple, independent routers. Each partition has its own Junos control plane, which runs as a virtual machine (VM), and a dedicated set of line cards. Each partition is called a guest network function (GNF).

The remaining statements are explained separately.

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

**Related Documentation**

- [Configuring an MX Series Router to Operate in BSYS Mode on page 27](#)
- [Configuring Guest Network Functions on page 36](#)

## guest-network-functions

---

```
Syntax  guest-network-functions {
        gnf id {
            control-plane-bandwidth-percent percent;
            description description;
            fpcs fpcs;
            af-name
        }
    }
```

**Hierarchy Level** [edit chassis network-slices]

**Release Information** Statement introduced in Junos OS Release 17.2.

**Description** Configure a guest network function for Junos Node Slicing.

Junos Node Slicing enables a single MX Series router to be partitioned to appear as multiple, independent routers. Each partition has its own Junos control plane, which runs as a virtual machine (VM), and a dedicated set of line cards. Each partition is called a guest network function (GNF).

The remaining statements are explained separately.

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

**Related Documentation**

- [Configuring Guest Network Functions on page 36](#)
- [network-slices on page 68](#)

## gnf

---

<b>Syntax</b>	<pre>gnf id {     control-plane-bandwidth-percent <i>percent</i>;     description <i>description</i>;     fpcs <i>fpcs</i>;     af-name }</pre>
<b>Hierarchy Level</b>	[edit chassis network-slices guest-network-functions]
<b>Release Information</b>	Statement introduced in Junos OS Release 17.2.
<b>Description</b>	<p>Define a GNF by assigning an ID to it.</p> <p>Junos Node Slicing enables a single MX Series router to be partitioned to appear as multiple, independent routers. Each partition has its own Junos control plane, which runs as a virtual machine (VM), and a dedicated set of line cards. Each partition is called a guest network function (GNF).</p>
<b>Options</b>	<p><b>id</b>—GNF ID</p> <p><b>Range:</b> 1–10</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>view-level—To view this statement in the configuration.</p> <p>control-level—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Guest Network Functions on page 36</a></li><li>• <a href="#">network-slices on page 68</a></li><li>• <a href="#">guest-network-functions on page 69</a></li></ul>

## control-plane-bandwidth-percent (Node Slicing)

---

<b>Syntax</b>	control-plane-bandwidth-percent <i>percent</i> ;
<b>Hierarchy Level</b>	[edit chassis network-slices guest-network-functions gnf]
<b>Release Information</b>	Statement introduced in Junos OS Release 17.2.
<b>Description</b>	Allocate a percentage of the bandwidth that exists on the control plane on the router to the specified guest network function (GNF). Allocating bandwidth prevents potential overutilization by one GNF over another.
<b>Options</b>	<i>percent</i> —Percentage of control plane bandwidth.
<b>Required Privilege Level</b>	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">network-slices on page 68</a></li><li>• <a href="#">guest-network-functions on page 69</a></li><li>• <a href="#">gnf on page 70</a></li><li>• <a href="#">description on page 72</a></li><li>• <a href="#">fpcs on page 73</a></li></ul>

## description (GNF)

---

<b>Syntax</b>	<code>description <i>description</i>;</code>
<b>Hierarchy Level</b>	[edit chassis network-slices guest-network-functions gnf]
<b>Release Information</b>	Statement introduced in Junos OS Release 17.2.
<b>Description</b>	Provide a description string for the specified guest network function (GNF).
<b>Options</b>	<b><i>description</i></b> —A description string for the specified guest network function (GNF).
<b>Required Privilege Level</b>	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">network-slices on page 68</a></li><li>• <a href="#">guest-network-functions on page 69</a></li><li>• <a href="#">gnf on page 70</a></li><li>• <a href="#">control-plane-bandwidth-percent on page 71</a></li><li>• <a href="#">fpcs on page 73</a></li></ul>



## fpcs (Node Slicing)

<b>Syntax</b>	<code>fpcs fpcs;</code>
<b>Hierarchy Level</b>	[edit chassis network-slices guest-network-functions gnf]
<b>Release Information</b>	Statement introduced in Junos OS Release 17.2.
<b>Description</b>	Assign FPCs to a guest network function.



**NOTE:** A given FPC cannot be assigned to more than one GNF.

<b>Options</b>	<i>fpcs</i> —The FPC to be assigned to a guest network function. Range: 0–63
<b>Required Privilege Level</b>	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">network-slices on page 68</a></li> <li>• <a href="#">guest-network-functions on page 69</a></li> <li>• <a href="#">gnf on page 70</a></li> <li>• <a href="#">control-plane-bandwidth-percent on page 71</a></li> <li>• <a href="#">description on page 72</a></li> </ul>

## af-name

---

**Syntax**

```
af-name {  
  peer-gnf {  
    id peer-gnf-id;  
    remote-af-name;  
  }  
  description af-description;  
}
```

**Hierarchy Level** [edit chassis network-slices guest-network-functions gnf id]

**Release Information** Statement introduced in Junos OS Release 17.4R1.

**Description** Configure Abstracted Fabric (AF) interface between a pair of guest network functions (GNFs). AF interface is a pseudo interface that represents a first class Ethernet interface behavior. An AF interface is created on a GNF to communicate with the peer GNF when the two GNFs are connected to each other through the CLI.



**NOTE:** Only one AF interface can be configured between a pair of GNFs.

---

**Options** *af-name*—Name of the AF interface being created on the GNF.  
**Range:** af0 through af9

*id peer-gnf-id*—Name of the GNF peer connected using the AF interface.  
**Range:** 1 through 10

*remote-af-name*—Name of the AF interface on the peer GNF.

*description af-description*—A description for the AF interface.

**Required Privilege Level** system

**Related Documentation**

- [Configuring Abstracted Fabric Interfaces Between a Pair of GNFs on page 41](#)

## peer-gnf

---

<b>Syntax</b>	<pre>peer-gnf {   id <i>peer-gnf-id</i>;   remote-af-name; }</pre>
<b>Hierarchy Level</b>	[edit chassis network-slices guest-network-functions gnf id <i>af-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 17.4R1.
<b>Description</b>	Configure the details of the GNF peer connected using the AF interface.
<b>Options</b>	<p><b>id</b><i>peer-gnf-id</i>—Name of the GNF peer connected using the Abstracted Fabric (AF) interface.</p> <p><b>Range:</b> 1 through 10</p> <p><b>remote-af-name</b>—Name of the AF interface on the peer GNF.</p>
<b>Required Privilege Level</b>	system
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Abstracted Fabric Interfaces Between a Pair of GNFs on page 41</a></li> </ul>

## description (AF)

---

<b>Syntax</b>	<code>description <i>description</i>;</code>
<b>Hierarchy Level</b>	[edit chassis network-slices guest-network-functions gnf af-name]
<b>Release Information</b>	Statement introduced in Junos OS Release 17.4R1.
<b>Description</b>	Provide a description string for the specified Abstracted Fabric (AF) interface.
<b>Options</b>	<b><i>description</i></b> —A description string for the specified Abstracted Fabric (AF) interface.
<b>Required Privilege Level</b>	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Abstracted Fabric Interfaces Between a Pair of GNFs on page 41</a></li><li>• <a href="#">gnf on page 70</a></li><li>• <a href="#">control-plane-bandwidth-percent on page 71</a></li><li>• <a href="#">fpcs on page 73</a></li></ul>

## CHAPTER 8

# Operational Commands for BSYS

- `show chassis network-slices`
- `show chassis fpc pic-status` (Node Slicing)
- `show chassis fpc` (Node Slicing)
- `show chassis adc` (Node Slicing)
- `show chassis network-slices fpcs`
- `show system anomalies gnf-id` (BSYS)

## show chassis network-slices

<b>Syntax</b>	<code>show chassis network-slices</code> <code>&lt;gnf <i>gnf-id</i>&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 17.2R1.
<b>Description</b>	Display Junos Node Slicing information for the guest network functions (GNFs) configured on the base system (BSYS). The <b>gnf <i>gnf-id</i></b> option displays the information about a particular GNF.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show chassis fpc pic-status (Node Slicing) on page 81</a></li> <li>• <a href="#">show chassis fpc (Node Slicing) on page 83</a></li> <li>• <a href="#">show chassis adc (Node Slicing) on page 85</a></li> <li>• <a href="#">show chassis network-slices fpcs on page 87</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show chassis network-slices on page 79</a> <a href="#">show chassis network-slices gnf on page 79</a> <a href="#">show chassis network-slices gnf &lt;gnf-id&gt; on page 80</a>
<b>Output Fields</b>	<a href="#">Table 8 on page 78</a> lists the output fields for the <b>show chassis network-slices</b> command. Output fields are listed in the approximate order in which they appear.

*Table 8: show chassis network-slices Output Fields*

Field Name	Field Description
GNF	GNF identifier for each partition.
Description	Description of the guest network function.
State	Status of the GNF. <ul style="list-style-type: none"> <li>• <b>Online</b>—GNF online and running.</li> <li>• <b>Offline</b>—GNF is offline.</li> </ul>
Uptime	Duration for which the GNFs have been up and running.
GNF ID	Shows the GNF ID.
GNF description	Shows the description of GNF.
FPCs assigned	Shows the FPC slot numbers assigned to the GNF.

*Table 8: show chassis network-slices Output Fields (continued)*

Field Name	Field Description
FPCs online	Shows the FPCs that are online.
BSYS	Shows the name of the BSYS.
BSYS sw version	Shows the Junos software version used in the BSYS.
GNF sw version	Shows the Junos software version used in the GNF.
BSYS master RE	Master Routing Engine slot.
GNF uptime	Duration for which the GNF has been up and running.
GNF Routing Engine Status:	Show the details of the Routing Engines in the slot 0 and 1. It includes the Current State, (master or backup), Routing Engine model, and GNF host name.

## Sample Output

### show chassis network-slices

```

user@router> show chassis network-slices
guest-network-functions:
GNF   Description   State      Uptime
  1   gnf-a           Online     12 hours, 46 minutes, 11 seconds
  2   gnf-b           Online     12 hours, 13 minutes, 57 seconds
  3   gnf-c           Online     12 hours, 3 minutes, 55 seconds
  4   gnf-d           Online     12 hours, 8 minutes, 20 seconds
  5   gnf-e           Online     12 hours, 2 minutes, 46 seconds
  6   gnf-f           Online     11 hours, 56 minutes, 29 seconds

```

### show chassis network-slices gnf

```

user@router> show chassis network-slices gnf

GNF ID          1
GNF description NA
GNF state       Online
FPCs assigned   7 8
FPCs online     7 8
BSYS            router(mx960)
BSYS sw version 18.2-20180321_0948_bsys
GNF sw version  18.2-20180314_gnf
Chassis         mx960
BSYS master RE  0
GNF uptime      4 days, 23 hours, 55 minutes, 1 second
GNF Routing Engine Status:
Slot 0:
  Current state Master
  Model         RE-GNF-2400x4
  GNF host name gnf-host0
Slot 1:
  Current state Backup

```

```
Model RE-GNF-2400x4
GNF host name gnf-host1
GNF ID 2
GNF description NA
GNF state Online
FPCs assigned NA
FPCs online NA
BSYS router(mx960)
BSYS sw version 18.2-20180321_0948_bsys
GNF sw version 18.2-20180216_gnf
Chassis mx960
BSYS master RE 0
GNF uptime 4 days, 23 hours, 53 minutes, 54 seconds
GNF Routing Engine Status:
Slot 0:
  Current state Master
  Model RE-GNF-2400x4
  GNF host name gnf-host2
Slot 1:
  Current state Backup
  Model RE-GNF-2400x4
  GNF host name gnf-host3
```

#### `show chassis network-slices gnf <gnf-id>`

```
user@router> show chassis network-slices gnf 2

GNF ID 2
GNF description NA
GNF state Online
FPCs assigned NA
FPCs online NA
BSYS router(mx960)
BSYS sw version 18.2-20180321_0948_bsys
GNF sw version 18.2-20180216_gnf
Chassis mx960
BSYS master RE 0
GNF uptime 4 days, 23 hours, 53 minutes, 54 seconds
GNF Routing Engine Status:
Slot 0:
  Current state Master
  Model RE-GNF-2400x4
  GNF host name gnf-host2
Slot 1:
  Current state Backup
  Model RE-GNF-2400x4
  GNF host name gnf-host3
```



## show chassis fpc pic-status (Node Slicing)

<b>Syntax</b>	show chassis fpc pic-status
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 Universal Routing Platforms.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 Universal Routing Platforms.</p>
<b>Description</b>	Display the status of the physical interface cards (PICs) of each Flexible PIC Concentrator (FPC) assigned to different guest network functions (GNFs).
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show chassis network-slices on page 78</a></li> <li>• <a href="#">show chassis fpc (Node Slicing) on page 83</a></li> <li>• <a href="#">show chassis adc (Node Slicing) on page 85</a></li> <li>• <a href="#">show chassis network-slices fpcs on page 87</a></li> </ul>

## Sample Output

```
user@router> show chassis fpc pic-status
```

```

Slot 0  Online      MPC5E 3D 24XGE+6XLGE      GNF 3
      PIC 0 Online      12X10GE SFPP OTN
      PIC 1 Offline     12X10GE SFPP OTN
      PIC 2 Offline     3X40GE QSFPP
      PIC 3 Online      3X40GE QSFPP
Slot 1  Online      MPC9E 3D      GNF 2
      PIC 1 Online      MRATE-12xQSFPP-XGE-XLGE-CGE
Slot 2  Online      MPC5E 3D Q 2CGE+4XGE      GNF 3
      PIC 0 Online      2X10GE SFPP OTN
      PIC 1 Online      1X100GE CFP2 OTN
      PIC 2 Online      2X10GE SFPP OTN
      PIC 3 Online      1X100GE CFP2 OTN
Slot 3  Online      MPCE Type 2 3D EQ      GNF 6
Slot 4  Online      MPC6E 3D      GNF 6
      PIC 0 Online      24X10GE SFPP
      PIC 1 Online      2X100GE CFP2 OTN
Slot 5  Online      MPC9E 3D      GNF 4
      PIC 0 Online      MRATE-12xQSFPP-XGE-XLGE-CGE
Slot 6  Online      MPC7E 3D MRATE-12xQSFPP-XGE-XLGE-CGE      GNF 1
      PIC 0 Online      MRATE-6xQSFPP-XGE-XLGE-CGE
      PIC 1 Online      MRATE-6xQSFPP-XGE-XLGE-CGE
Slot 7  Online      MPC5E 3D 2CGE+4XGE      GNF 5
      PIC 0 Online      2X10GE SFPP OTN
      PIC 1 Online      1X100GE CFP2 OTN
      PIC 2 Online      2X10GE SFPP OTN
      PIC 3 Online      1X100GE CFP2 OTN
Slot 8  Online      MPC6E 3D      GNF 5
      PIC 0 Online      24X10GE SFPP OTN

```

Slot 9	Online	MPC6E 3D	GNF 5
PIC 0	Online	24X10GE SFPP	
PIC 1	Online	4X100GE CXP	

## show chassis fpc (Node Slicing)

**Syntax** `show chassis fpc`

**Release Information** Command introduced before Junos OS Release 7.4.  
Command introduced in Junos OS Release 12.3 for MX2020 Universal Routing Platforms.  
Command introduced in Junos OS Release 12.3 for MX2010 Universal Routing Platforms.

**Description** Display information about Flexible PIC Concentrators (fpcs) assigned to different guest network functions (GNFs).

**Required Privilege Level** view

**Related Documentation**

- [show chassis network-slices on page 78](#)
- [show chassis fpc pic-status \(Node Slicing\) on page 81](#)
- [show chassis adc \(Node Slicing\) on page 85](#)
- [show chassis network-slices fpcs on page 87](#)

**Output Fields** [Table 9 on page 83](#) lists the output fields for the `show chassis fpc` command. Output fields are listed in the approximate order in which they appear.

*Table 9: show chassis fpc Output Fields*

Field Name	Field Description
<b>Slot or Slot State</b>	Slot number and state. The state can be one of the following conditions: <ul style="list-style-type: none"> <li>• <b>Dead</b>—Held in reset because of errors.</li> <li>• <b>Diag</b>—Slot is being ignored while the FPC is running diagnostics.</li> <li>• <b>Dormant</b>—Held in reset.</li> <li>• <b>Empty</b>—No FPC is present.</li> <li>• <b>Online</b>—FPC is online and running.</li> <li>• <b>Present</b>—FPC is detected by the chassis daemon but either is not supported by the current version of Junos OS or is inserted in the wrong slot. The output also states either <b>Hardware Not Supported</b> or <b>Hardware Not In Right Slot</b>. The FPC is coming up but not yet online.</li> <li>• <b>Probed</b>—Probe is complete; awaiting restart of the Packet Forwarding Engine.</li> <li>• <b>Probe-wait</b>—Waiting to be probed.</li> </ul>
<b>Temp (C) or Temperature</b>	Temperature of the air passing by the FPC, in degrees Celsius or in both Celsius and Fahrenheit.
<b>Total CPU Utilization (%)</b>	Total percentage of CPU being used by the FPC's processor.

Table 9: show chassis fpc Output Fields (continued)

Field Name	Field Description
<b>Interrupt CPU Utilization (%)</b>	Of the total CPU being used by the FPC's processor, the percentage being used for interrupts.
<b>1 min CPU Utilization (%)</b>	Information about the Routing Engine's CPU utilization in the past 1 minute.
<b>5 min CPU Utilization (%)</b>	Information about the Routing Engine's CPU utilization in the past 5 minutes.
<b>15 min CPU Utilization (%)</b>	Information about the Routing Engine's CPU utilization in the past 15 minutes.
<b>Memory DRAM (MB)</b>	Total DRAM, in megabytes, available to the FPC's processor.
<b>Heap Utilization (%)</b>	Percentage of heap space (dynamic memory) being used by the FPC's processor. If this number exceeds 80 percent, there may be a software problem (memory leak).  <b>NOTE:</b> On MX Series routers and EX Series switches in a broadband edge environment, heap utilization levels higher than 70 percent can affect unified ISSU, router stability, or scaling capability.
<b>Buffer Utilization (%)</b>	Percentage of buffer space being used by the FPC's processor for buffering internal messages.
<b>GNF</b>	GNF identifier associated with each FPC.  (pic-status output only) GNF identifier associated with each PIC.

## Sample Output

```
user@router> show chassis fpc
```

Slot	State	Temp (C)	CPU Utilization (%)		CPU Utilization(%)			Memory Utilization (%)				GNF
			Total	Interrupt	1min	5min	15min	DRAM(MB)	Heap	Buffer		
0	Online	45	12	0	12	12	12	3584	6	25	3	
1	Online	57	22	0	20	20	20	3136	16	22	2	
2	Online	50	19	0	17	17	16	3584	6	25	3	
3	Online	28	10	0	11	11	11	2048	10	20	6	
4	Online	42	20	0	20	19	19	3584	8	25	6	
5	Online	58	22	0	21	20	20	3136	16	22	4	
6	Online	49	17	0	15	16	16	3136	13	20	1	
7	Online	44	11	0	10	10	10	3584	6	25	5	
8	Online	40	19	0	18	18	18	3584	8	25	5	
9	Online	44	19	0	20	20	20	3584	8	25	5	

## show chassis adc (Node Slicing)

<b>Syntax</b>	show chassis adc
<b>Release Information</b>	Command introduced in Junos OS Release 12.3 for MX2020 Universal Routing Platforms. Command introduced in Junos OS Release 12.3 for MX2010 Universal Routing Platforms. Command introduced in Junos OS Release 17.2 for MX2008 Universal Routing Platforms.
<b>Description</b>	Display chassis information about the adapter cards (ADCs).
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show chassis network-slices on page 78</a></li> <li>• <a href="#">show chassis fpc pic-status (Node Slicing) on page 81</a></li> <li>• <a href="#">show chassis fpc (Node Slicing) on page 83</a></li> <li>• <a href="#">show chassis network-slices fpcs on page 87</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show chassis adc (Node Slicing) on page 85</a>
<b>Output Fields</b>	<a href="#">Table 10 on page 85</a> lists the output fields for the <b>show chassis adc</b> command. Output fields are listed in the approximate order in which they appear.

Table 10: show chassis adc Output Fields

Field Name	Field Description
Slot	Slot number.
State	Status of the adapter card. <ul style="list-style-type: none"> <li>• <b>Online</b>—The adapter card is online and running.</li> <li>• <b>Offline</b>—Adapter card is powered down.</li> </ul>
Uptime	How long the Routing Engine has been connected to the adapter card and, therefore, how long the adapter card has been up and running.
GNF (Node slicing)	GNF identifier for each ADC.

## Sample Output

### show chassis adc (Node Slicing)

```

user@router> show chassis adc
Slot  State          Uptime                               GNF
0     Online           12 hours, 57 minutes, 46 seconds    3
1     Empty          --- Native line card ---          2

```

2	Online	12 hours, 57 minutes, 18 seconds	3
3	Online	11 minutes, 23 seconds	6
4	Empty	--- Native line card ---	6
5	Empty	--- Native line card ---	4
6	Online	13 hours, 38 minutes, 58 seconds	1
7	Online	13 hours, 3 minutes, 40 seconds	5
8	Empty	--- Native line card ---	5
9	Empty	--- Native line card ---	5

## show chassis network-slices fpcs

<b>Syntax</b>	show chassis network-slices fpcs
<b>Release Information</b>	Command introduced in Junos OS Release 17.2R1.
<b>Description</b>	Display information about the FPCs associated with different guest network functions (GNFs).
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show chassis network-slices on page 78</a></li> <li>• <a href="#">show chassis fpc pic-status (Node Slicing) on page 81</a></li> <li>• <a href="#">show chassis fpc (Node Slicing) on page 83</a></li> <li>• <a href="#">show chassis adc (Node Slicing) on page 85</a></li> </ul>
<b>Output Fields</b>	Table 11 on page 87 lists the output fields for the <b>show chassis network-slices fpcs</b> command. Output fields are listed in the approximate order in which they appear.

Table 11: show chassis network-slices fpcs Output Fields

Field Name	Field Description
GNF	GNF ID.
FPCs	Slot numbers of FPCs associated with the GNF ID.

## Sample Output

```

user@router> show chassis network-slices fpcs
guest-network-functions:
GNF      FPCs
1         6
2         1
3         0 2
4         5
5         7 8 9
6         3 4

```

## show system anomalies gnf-id (BSYS)

<b>Syntax</b>	<b>show system anomalies gnf-id <i>id</i> (all-anomalies   config   fru   system)</b>
<b>Release Information</b>	Command introduced in Junos OS Release 17.4R1.
<b>Description</b>	Display incompatibilities between the software version running on the base system (BSYS) and the software running on a specific guest network function (GNF).
<b>Options</b>	<p><b>gnf-id <i>id</i></b>—Specify the GNF ID for which you want to view the software incompatibilities.</p> <p><b>all-anomalies</b>—Display the multi-version software incompatibilities from all categories—system, configuration and FRU.</p> <p><b>config</b>—Display feature-level incompatibilities between software versions.</p> <p><b>fru</b>—Display the FRU-level incompatibilities between software versions. This can be an incompatibility pertaining to the support for a specific FRU.</p> <p><b>system</b>—Display the system-level incompatibilities between software versions. These include IPC incompatibility, CLI or SNMP incompatibility.</p>
<b>Required Privilege Level</b>	View
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show system anomalies (GNF) on page 112</a></li> <li>• <a href="#">Generic Guidelines for Using JDM Server Commands on page 123</a></li> <li>• <a href="#">request virtual-network-functions on page 131</a></li> <li>• <a href="#">request server authenticate-peer-server on page 130</a></li> </ul>
<b>Output Fields</b>	<a href="#">Table 12 on page 88</a> lists the output fields for the <b>show system anomalies gnf-id</b> command. Output fields are listed in the approximate order in which they appear.

*Table 12: show system anomalies gnf-id Output Fields*

Field Name	Field Description
<b>Anomaly Type</b>	Shows the software incompatibility type. The following are the possible values: <ul style="list-style-type: none"> <li>• SYS—Indicates system-specific incompatibilities.</li> <li>• FRU—Indicates FRU-specific incompatibilities.</li> <li>• CONFIG—Indicates feature-specific incompatibilities.</li> </ul>
<b>Anomaly ID</b>	Shows the incompatibility ID.



Table 12: show system anomalies gnf-id Output Fields (continued)

Field Name	Field Description
<b>Default Action</b>	Shows the default actions associated with incompatibilities. The following are the possible values: <ul style="list-style-type: none"> <li>WARN—Indicates the presence of a minor incompatibility. It causes a warning.</li> <li>ABORT—Indicates the presence of a major incompatibility. It causes an abort action.</li> </ul>
<b>Class</b>	Indicates of the incompatibility is present in GNF, or BSYS, or both.
<b>Message</b>	Incompatibility description.
<b>FRU ID</b>	Field-replaceable unit (FRU) ID. Applicable in the case of FRU-specific incompatibilities.
<b>Firmware</b>	Firmware of the FRU.

## Sample Output

```
user@router> show system anomalies gnf-id 1
```

Anomaly Type	Anomaly ID FRU ID	Default Action Firmware	Class	Message
SYS incompatibility 1	100	WARN	GNF Present	system
SYS incompatibility 2	103	WARN	BOTH Present	system
SYS incompatibility 3	200	WARN	BSYS Present	system



## CHAPTER 9

# Operational Commands for GNF

- `show chassis gnf`
- `show chassis gnf`
- `show chassis hardware (GNF)`
- `show chassis fpc (GNF)`
- `show chassis fpc pic-status (GNF)`
- `show chassis adc (GNF)`
- `show interfaces (Abstracted Fabric)`
- `show system anomalies (GNF)`

## show chassis gnf

**Syntax** `show chassis gnf`

**Release Information** Command introduced in Junos OS Release 18.2R1.

**Description** Display information about the guest network function (GNF) you logged in.

**Required Privilege Level** view

**Related Documentation**

- [show chassis hardware \(GNF\) on page 96](#)
- [show chassis fpc pic-status \(GNF\) on page 100](#)
- [show chassis adc \(GNF\) on page 101](#)
- [show chassis network-slices on page 78](#)

**List of Sample Output** [show chassis gnf on page 93](#)

**Output Fields** [Table 13 on page 92](#) lists the output fields for the **show chassis gnf** command. Output fields are listed in the approximate order in which they appear.

*Table 13: show chassis gnf Output Fields*

Field Name	Field Description
GNF ID	Shows the GNF ID.
GNF State	Status of the GNF. <ul style="list-style-type: none"> <li>• <b>Online</b>—GNF online and running.</li> <li>• <b>Offline</b>—GNF is offline.</li> </ul>
GNF description	Description of the guest network function.
FPCs assigned	The FPC slot numbers assigned to the GNF.
FPCs online	The FPCs that are online.
BSYS	Name of the BSYS.
BSYS sw version	The Junos software version used in the BSYS.
GNF sw version	The Junos software version used in the GNF.
BSYS master RE	Master Routing Engine slot.
GNF uptime	Duration for which the GNF has been up and running.

Table 13: show chassis gnf Output Fields (continued)

Field Name	Field Description
GNF Routing Engine Status:	The details of the Routing Engines in the slot 0 and 1. The details include the Current State, (master or backup), Routing Engine model, and GNF host name.

## Sample Output

### show chassis gnf

```

user@router> show chassis gnf

GNF ID          1
GNF description  NA
GNF state       Online
FPCs assigned   8 9
FPCs online     8 9
BSYS            router(mx960)
BSYS sw version 18.2-20180321_0948_bsys
GNF sw version  18.2-20180314_1035_gnf
Chassis         mx960
BSYS master RE  0
GNF uptime      54 minutes, 37 seconds
GNF Routing Engine Status:
Slot 0:
  Current state  Master
  Model         RE-GNF-2100x4
  GNF host name  gnf-host0
Slot 1:
  Current state  Backup
  Model         RE-GNF-2100x4
  GNF host name  gnf-host1

```

## show chassis gnf

**Syntax** `show chassis gnf`

**Release Information** Command introduced in Junos OS Release 18.2R1.

**Description** Display information about the guest network function (GNF) you logged in.

**Required Privilege Level** view

**Related Documentation**

- [show chassis hardware \(GNF\) on page 96](#)
- [show chassis fpc pic-status \(GNF\) on page 100](#)
- [show chassis adc \(GNF\) on page 101](#)
- [show chassis network-slices on page 78](#)

**List of Sample Output** [show chassis gnf on page 95](#)

**Output Fields** [Table 13 on page 92](#) lists the output fields for the **show chassis gnf** command. Output fields are listed in the approximate order in which they appear.

*Table 14: show chassis gnf Output Fields*

Field Name	Field Description
GNF ID	Shows the GNF ID.
GNF State	Status of the GNF. <ul style="list-style-type: none"> <li>• <b>Online</b>—GNF online and running.</li> <li>• <b>Offline</b>—GNF is offline.</li> </ul>
GNF description	Description of the guest network function.
FPCs assigned	The FPC slot numbers assigned to the GNF.
FPCs online	The FPCs that are online.
BSYS	Name of the BSYS.
BSYS sw version	The Junos software version used in the BSYS.
GNF sw version	The Junos software version used in the GNF.
BSYS master RE	Master Routing Engine slot.
GNF uptime	Duration for which the GNF has been up and running.

*Table 14: show chassis gnf Output Fields (continued)*

Field Name	Field Description
GNF Routing Engine Status:	The details of the Routing Engines in the slot 0 and 1. The details include the Current State, (master or backup), Routing Engine model, and GNF host name.

## Sample Output

### show chassis gnf

```

user@router> show chassis gnf

GNF ID          1
GNF description  NA
GNF state       Online
FPCs assigned   8 9
FPCs online     8 9
BSYS            router(mx960)
BSYS sw version 18.2-20180321_0948_bsys
GNF sw version  18.2-20180314_1035_gnf
Chassis         mx960
BSYS master RE  0
GNF uptime      54 minutes, 37 seconds
GNF Routing Engine Status:
Slot 0:
  Current state  Master
  Model         RE-GNF-2100x4
  GNF host name  gnf-host0
Slot 1:
  Current state  Backup
  Model         RE-GNF-2100x4
  GNF host name  gnf-host1

```

## show chassis hardware (GNF)

**Syntax** `show chassis hardware`

**Release Information** Command introduced before Junos OS Release 7.4.  
Command introduced in Junos OS Release 12.3 for MX2010 and MX2020 Universal Routing Platforms.

**Description** Display a list of all hardware components of the chassis, including the hardware version level and serial number, the GNF Routing Engine details, and the FPCs assigned to the GNF.

**Required Privilege Level** view

**Related Documentation**

- [show chassis fpc \(GNF\) on page 98](#)
- [show chassis fpc pic-status \(GNF\) on page 100](#)
- [show chassis adc \(GNF\) on page 101](#)

**Output Fields** [Table 15 on page 96](#) lists the output fields for the **show chassis hardware** command. Output fields are listed in the approximate order in which they appear.

*Table 15: show chassis hardware Output Fields*

Field Name	Field Description
<b>Item</b>	Chassis component: <ul style="list-style-type: none"> <li>• Information about the backplane, Routing Engine, Power Entry Modules (PEMs), and fan trays. Also displays information about Flexible PIC Concentrators (FPCs) and associated Physical Interface Cards (PICs), Modular Port Concentrators (MPCs) and associated Modular Interface Cards (MICs), or Dense Port Concentrators (DPCs).</li> </ul>
<b>Version</b>	Revision level of the chassis component.
<b>Part number</b>	Part number of the chassis component.
<b>Serial number</b>	Serial number of the chassis component. The serial number of the backplane is also the serial number of the router chassis. Use this serial number when you need to contact Juniper Networks Customer Support about the router or switch chassis.
<b>Description</b>	Brief description of the hardware item.



## Sample Output

```
user@router> show chassis hardware
```

```
bsys-re0:
```

```
-----
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN11C9CDDAFK	MX2010
Midplane	REV 35	750-044636	ABAB9184	Lower Backplane
Midplane 1	REV 02	711-044557	ABAB9048	Upper Backplane
PMP	REV 04	711-032426	ACAJ2622	Power Midplane
FPM Board	REV 09	760-044634	ABCF2618	Front Panel Display
PSM 0	REV 01	740-050037	1EDB3130084	DC 52V Power Supply
Module				
PSM 1	REV 01	740-050037	1EDB313001Z	DC 52V Power Supply
Module				
PSM 2	REV 01	740-050037	1EDB321018D	DC 52V Power Supply
Module				
PSM 3	REV 01	740-050037	1EDB32101AZ	DC 52V Power Supply
Module				
PSM 4	REV 01	740-050037	1EDB32202C2	DC 52V Power Supply
Module				
PSM 5	REV 01	740-050037	1EDB32100TC	DC 52V Power Supply
Module				
PSM 6	REV 01	740-050037	1EDB3210166	DC 52V Power Supply
Module				
PSM 7	REV 01	740-050037	1EDB3210165	DC 52V Power Supply
Module				
PSM 8	REV 01	740-050037	1EDB3210163	DC 52V Power Supply
Module				
PDM 0	REV 03	740-045234	1EGA3170177	DC Power Dist Module
Routing Engine 0	REV 08	750-055814	CAFN5537	RE-S-2X00x8
CB 0	REV 08	750-055087	CAFN3426	MX2K Enhanced SCB
Xcvr 0	REV 01	740-031980	ALMOHC7	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	123363A00418	SFP+-10G-SR
CB 1	REV 08	750-055087	CAFN3423	MX2K Enhanced SCB
SPMB 0	REV 05	711-041855	CAEZ5998	PMB Board
SPMB 1	REV 05	711-041855	CAEZ5993	PMB Board
SFB 0	REV 06	711-044466	ABCD6742	Switch Fabric Board
SFB 1	REV 06	711-044466	ABCG5627	Switch Fabric Board
SFB 2	REV 06	711-044466	ABCG5659	Switch Fabric Board
SFB 3	REV 06	711-044466	ABCG5653	Switch Fabric Board
SFB 4	REV 06	711-044466	ABCG5611	Switch Fabric Board
SFB 5	REV 06	711-044466	ABCG5635	Switch Fabric Board
SFB 6	REV 06	711-044466	ABCG5638	Switch Fabric Board
SFB 7	REV 06	711-044466	ABCG3650	Switch Fabric Board
FPC 8	REV 68	750-044130	ABCY5967	MPC6E 3D
CPU	REV 12	711-045719	ABCY9696	RMPC PMB
Fan Tray 0	REV 06	760-046960	ACAY0428	172mm FanTray - 6 Fans
Fan Tray 1	REV 06	760-046960	ACAY0800	172mm FanTray - 6 Fans
Fan Tray 2	REV 06	760-046960	ACAY0797	172mm FanTray - 6 Fans
Fan Tray 3	REV 06	760-046960	ACAY1047	172mm FanTray - 6 Fans

```
gnf2-re0:
```

```
-----
```

Chassis	GN59081553B0	MX2010-GNF	<<<
Routing Engine 0		RE-GNF-1700x4	

## show chassis fpc (GNF)

**Syntax** `show chassis fpc`

**Release Information** Command introduced before Junos OS Release 7.4.  
Command introduced in Junos OS Release 12.3 for MX2020 Universal Routing Platforms.  
Command introduced in Junos OS Release 12.3 for MX2010 Universal Routing Platforms.

**Description** Display information about the Flexible PIC Concentrators (fpcs) assigned to the guest network function (GNF).

**Required Privilege Level** view

**Related Documentation**

- [show chassis hardware \(GNF\) on page 96](#)
- [show chassis fpc pic-status \(GNF\) on page 100](#)
- [show chassis adc \(GNF\) on page 101](#)

**Output Fields** [Table 16 on page 98](#) lists the output fields for the `show chassis fpc` command. Output fields are listed in the approximate order in which they appear.

*Table 16: show chassis fpc Output Fields*

Field Name	Field Description
Slot or Slot State	Slot number and state. The state can be one of the following conditions: <ul style="list-style-type: none"> <li>• <b>Dead</b>—Held in reset because of errors.</li> <li>• <b>Diag</b>—Slot is being ignored while the FPC is running diagnostics.</li> <li>• <b>Dormant</b>—Held in reset.</li> <li>• <b>Empty</b>—No FPC is present.</li> <li>• <b>Online</b>—FPC is online and running.</li> <li>• <b>Present</b>—FPC is detected by the chassis daemon but either is not supported by the current version of Junos OS or is inserted in the wrong slot. The output also states either <b>Hardware Not Supported</b> or <b>Hardware Not In Right Slot</b>. The FPC is coming up but not yet online.</li> <li>• <b>Probed</b>—Probe is complete; awaiting restart of the Packet Forwarding Engine.</li> <li>• <b>Probe-wait</b>—Waiting to be probed.</li> </ul>
Temp (C) or Temperature	Temperature of the air passing by the FPC, in degrees Celsius or in both Celsius and Fahrenheit.
Total CPU Utilization (%)	Total percentage of CPU being used by the FPC's processor.
Interrupt CPU Utilization (%)	Of the total CPU being used by the FPC's processor, the percentage being used for interrupts.

Table 16: show chassis fpc Output Fields (continued)

Field Name	Field Description
1 min CPU utilization (%)	Information about the Routing Engine's CPU utilization in the past 1 minute.
5 min CPU utilization (%)	Information about the Routing Engine's CPU utilization in the past 5 minutes.
15 min CPU utilization (%)	Information about the Routing Engine's CPU utilization in the past 15 minutes.
Memory DRAM (MB)	Total DRAM, in megabytes, available to the FPC's processor.
Heap Utilization (%)	<p>Percentage of heap space (dynamic memory) being used by the FPC's processor. If this number exceeds 80 percent, there might be a software problem (memory leak).</p> <p><b>NOTE:</b> On MX Series routers and EX Series switches in a broadband edge environment, heap utilization levels higher than 70 percent can affect unified ISSU, router stability, or scaling capability.</p>
Buffer Utilization (%)	Percentage of buffer space being used by the FPC's processor for buffering internal messages.
GNF	<p>GNF identifier associated with each FPC.</p> <p>(pic-status output only) GNF identifier associated with each PIC.</p>

## Sample Output

```

user@router>show chassis fpc
          Temp CPU Utilization(%) CPU Utilization(%) Memory Utilization(%)
Slot State  (C)  Total Interrupt   1min 5min 15min  DRAM(MB)  Heap  Buffer  GNF
-----
  4  Online  42    20         0    19   19   19   3584     8   25    2
  6  Online  46    12         0    11   11   11   3136     8   19    2

```

## show chassis fpc pic-status (GNF)

---

<b>Syntax</b>	show chassis fpc pic-status
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 12.3 for MX2020 Universal Routing Platforms. Command introduced in Junos OS Release 12.3 for MX2010 Universal Routing Platforms.
<b>Description</b>	Display the status of the physical interface cards (PICs) of each Flexible PIC Concentrator (FPC) assigned to the guest network function (GNF).
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show chassis hardware (GNF) on page 96</a></li><li>• <a href="#">show chassis fpc (GNF) on page 98</a></li><li>• <a href="#">show chassis adc (GNF) on page 101</a></li></ul>

## Sample Output

```
user@router> show chassis fpc pic-status

Slot 4  Online      MPC6E 3D                GNF 2
  PIC 0  Online      24X10GE SFPP
  PIC 1  Online      4X100GE CXP
Slot 6  Online      MPC7E 3D MRATE-12xQSFP-XGE-XLGE-CGE  GNF 2
  PIC 0  Online      MRATE-6xQSFP-XGE-XLGE-CGE
  PIC 1  Online      MRATE-6xQSFP-XGE-XLGE-CGE
```

## show chassis adc (GNF)

<b>Syntax</b>	show chassis adc
<b>Release Information</b>	Command introduced in Junos OS Release 12.3 for MX2020 Universal Routing Platforms. Command introduced in Junos OS Release 12.3 for MX2010 Universal Routing Platforms. Command introduced in Junos OS Release 17.2 for MX2008 Universal Routing Platforms.
<b>Description</b>	Display chassis information about the adapter cards (ADCs) assigned to the guest network function (GNF).
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show chassis hardware (GNF) on page 96</a></li> <li>• <a href="#">show chassis fpc (GNF) on page 98</a></li> <li>• <a href="#">show chassis fpc pic-status (GNF) on page 100</a></li> </ul>
<b>Output Fields</b>	<a href="#">Table 17 on page 101</a> lists the output fields for the <b>show chassis adc</b> command. The output fields are listed in the approximate order in which they appear.

*Table 17: show chassis adc Output Fields*

Field Name	Field Description
<b>Slot</b>	Slot number.
<b>State</b>	Status of the adapter card. <ul style="list-style-type: none"> <li>• <b>Online</b>—The adapter card is online and running.</li> <li>• <b>Offline</b>—Adapter card is powered down.</li> </ul>
<b>Uptime</b>	How long the Routing Engine has been connected to the adapter card and, therefore, how long the adapter card has been up and running.
<b>GNF (Node slicing)</b>	GNF identifier for each ADC.

## Sample Output

```

user@router> show chassis adc

Slot  State              Uptime                      GNF
 4    Empty              --- Native line card ---    2
 6    Online              50 minutes, 2 seconds      2

```

## show interfaces (Abstracted Fabric)

**Syntax** `show interfaces af-interface-name`  
`<brief | detail | extensive | terse>`  
`<descriptions>`  
`<media>`  
`<snmp-index snmp-index>`  
`<statistics>`

**Release Information** Command introduced before Junos OS Release 17.4R1.

**Description** Display status information for the specified Abstracted Fabric (AF) interface.



**NOTE:** Starting in Junos OS Release 18.3R1, the `show interfaces af-interface-name` output provides transmit statistics of each Packet Forwarding Engine peer list on a given AF interface, in addition to the physical interface statistics.

**Options** `brief | detail | extensive | terse`—(Optional) Display the specified level of output.

`descriptions`—(Optional) Display interface description strings.

`media`—(Optional) Display media-specific information about network interfaces.

`snmp-index snmp-index`—(Optional) Display information for the specified SNMP index of the interface.

`statistics`—(Optional) Display static interface statistics.

**Required Privilege Level** view

**Related Documentation**

- [Configuring Abstracted Fabric Interfaces Between a Pair of GNFs on page 41](#)

**List of Sample Output**

- [show interfaces af-interface-name on page 106](#)
- [show interfaces af-interface-name on page 106](#)
- [show interfaces af-interface-name brief on page 107](#)
- [show interfaces af-interface-name extensive on page 108](#)
- [show interfaces af-interface-name detail on page 109](#)
- [show interfaces af-interface-name statistics on page 110](#)

**Output Fields** [Table 18 on page 103](#) describes the output fields for the `show interfaces` (Abstracted Fabric) command. Output fields are listed in the approximate order in which they appear.

Table 18: show interfaces (Abstracted Fabric) Output Fields

Field Name	Field Description	Level of Output
<b>Physical Interface</b>		
<b>Physical interface</b>	Name and status of the physical interface.	All levels
<b>Interface index</b>	Index number of the physical interface, which reflects its initialization sequence.	<b>detail extensive</b> none
<b>SNMP ifIndex</b>	SNMP index number for the physical interface.	<b>detail extensive</b> none
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Type</b>	Interface type.	
<b>Link-level type</b>	Encapsulation being used on the physical interface.	All levels
<b>MTU</b>	Maximum transmission unit size on the physical interface.	All levels
<b>Speed</b>	Speed at which the interface is running.	All levels
<b>Device flags</b>	Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .	All levels
<b>Interface flags</b>	Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .	All levels
<b>Link type</b>	Link type. For example, Full-Duplex.	All levels
<b>Damping</b>	Damping information.	<b>detail extensive</b>
<b>Alternate link address</b>	Address of the alternative link.	<b>detail extensive</b>
<b>Hold-times</b>	Current interface hold-time up and hold-time down, in milliseconds (ms).	<b>detail extensive</b>
<b>Current address</b>	Configured MAC address.	<b>detail extensive</b> none
<b>Hardware address</b>	Hardware MAC address.	<b>detail extensive</b> none
<b>Last flapped</b>	Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> .	<b>detail extensive</b> none
<b>Statistics last cleared</b>	Time when the statistics for the interface were last set to zero.	<b>detail extensive</b>

Table 18: show interfaces (Abstracted Fabric) Output Fields (continued)

Field Name	Field Description	Level of Output
<b>Traffic statistics</b>	Number and rate of bytes and packets received and transmitted on the physical interface. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>	<b>detail extensive</b>
<b>IPv6 transit statistics</b>	Number of IPv6 transit bytes and packets received and transmitted on the interface if IPv6 statistics tracking is enabled.	<b>extensive</b>
<b>Input errors</b>	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>	<b>extensive</b>
<b>Output errors</b>	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue.</li> </ul> <p><b>NOTE:</b></p> <ul style="list-style-type: none"> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>	<b>extensive</b>
<b>Bandwidth</b>	Shows the interface bandwidth.	<b>detail extensive none</b>
<b>Peer GNF id</b>	The GNF peer connected using the AF interface.	<b>detail extensive none</b>
<b>Peer GNF Forwarding element(FE) view</b>	Shows forwarding element (FE) number and the FPC slot, FE bandwidth, and FE status (up/down).	<b>detail extensive none</b>
<b>Transmit Packets</b>	Shows the packets transmitted.	
<b>Transmit Bytes</b>	Shows the bytes transmitted.	
<b>Residual Transmit Statistics</b>	Displays the historical transmit statistics for the peer AF interface from the FPC events (FPC offline/online/removal/power off on the local or remote side).	



Table 18: show interfaces (Abstracted Fabric) Output Fields (continued)

Field Name	Field Description	Level of Output
Fabric Queue Statistics	Displays the statistics of the transmit traffic on the fabric queues (high and low queues) for each peer PFE on the AF interface.	
Residual Queue Statistics	Displays the historical fabric queue statistics for the peer AF interface from the FPC events (FPC offline/online/removal/power off on the local or remote side).	
<b>Logical Interface</b>		
<b>Logical interface</b>	Name of the logical interface.	All levels
<b>Index</b>	Index number of the logical interface, which reflects its initialization sequence.	<b>detail extensive</b> none
<b>SNMP ifIndex</b>	SNMP interface index number for the logical interface.	<b>detail extensive</b> none
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Flags</b>	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> .	All levels
<b>Encapsulation</b>	Encapsulation on the logical interface.	All levels
<b>Protocol</b>	Protocol family. Possible values are described in the “Protocol Field” section under <i>Common Output Fields Description</i> .	<b>detail extensive</b> none
<b>MTU</b>	Maximum transmission unit size on the logical interface.	<b>detail extensive</b> none
<b>Traffic statistics</b>	<p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> <li><b>Input bytes, Output bytes</b>—Number of bytes received and transmitted on the interface set.</li> <li><b>Input packets, Output packets</b>—Number of packets received and transmitted on the interface set.</li> </ul>	<b>detail extensive</b>
<b>Transit statistics</b>	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	<b>extensive</b>
<b>Local statistics</b>	Number and rate of bytes and packets destined to the router.	<b>extensive</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Route Table</b>	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	<b>detail extensive</b> none
<b>Flags</b>	Information about protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive</b>
<b>Addresses, Flags</b>	Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive</b> none

Table 18: show interfaces (Abstracted Fabric) Output Fields (continued)

Field Name	Field Description	Level of Output
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.	<b>brief</b>
<b>Flags</b>	Information about the address flag. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive none</b>
<b>Destination</b>	IP address of the remote side of the connection.	<b>detail extensive none</b>
<b>Local</b>	IP address of the logical interface.	<b>detail extensive none</b>
<b>Broadcast</b>	Broadcast address of the logical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>

## Sample Output

### show interfaces af-interface-name

```
user@router> show interfaces af4
```

```
Physical interface: af4, Enabled, Physical link is Up
  Interface index: 139, SNMP ifIndex: 514
  Type: Ethernet, Link-level type: Ethernet, MTU: 1518, Speed: 480000mbps
  Device flags   : Present Running
  Interface flags: Internal: 0x4000
  Link type      : Full-Duplex
  Current address: 2c:6b:f5:55:eb:f6, Hardware address: 2c:6b:f5:55:eb:f6
  Last flapped   : 2017-10-18 19:40:00 EDT (02:47:02 ago)
  Input rate      : 0 bps (0 pps)
  Output rate     : 0 bps (0 pps)
  Bandwidth       : 480 Gbps
  Peer GNF id     : 4
  Peer GNF Forwarding element(FE) view :
      FPC slot:FE Num      FE Bandwidth(Gbps)      Status
      6:0                  240                  Up
      6:1                  240                  Up

Logical interface af4.1 (Index 328) (SNMP ifIndex 593)
  Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ] Encapsulation: ENET2
  Input packets : 9
  Output packets: 9
  Protocol inet, MTU: 1500
  Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1,
  Curr new hold cnt: 0, NH drop cnt: 0
  Flags: Sendbroadcast-pkt-to-re
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255
  Protocol multiservice, MTU: Unlimited
```

### show interfaces af-interface-name

This output is applicable to Junos OS Release 18.3R1 onwards.

```

user@router> show interfaces af9
Physical interface: af9, Enabled, Physical link is Up
  Interface index: 209, SNMP ifIndex: 527
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 370000mbps
  Device flags   : Present Running
  Interface flags: Internal: 0x4000
  Link type      : Full-Duplex
  Link flags     : None
  Current address: 00:90:69:2b:00:4c, Hardware address: 00:90:69:2b:00:4c
  Last flapped   : 2018-09-12 01:44:01 PDT (00:01:02 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Bandwidth      : 370 Gbps
  Peer GNF id    : 9
  Peer GNF Forwarding element(FE) view :
  FPC slot:FE num  FE Bandwidth(Gbps) Status    Transmit Packets    Transmit
  Bytes
      6:0                130        Up              0
      0
     12:0                120        Up              0
      0
     12:1                120        Up              0
      0

  Residual Transmit Statistics :
  Packets :                0 Bytes :                0

  Fabric Queue Statistics :
  FPC slot:FE num  High priority(pkts)    Low priority(pkts)
      6:0                0                0
     12:0                0                0
     12:1                0                0
  FPC slot:FE num  High priority(bytes)    Low priority(bytes)
      6:0                0                0
     12:0                0                0
     12:1                0                0

  Residual Queue Statistics :
      High priority(pkts)    Low priority(pkts)
          0                0
      High priority(bytes)    Low priority(bytes)
          0                0

  Logical interface af9.0 (Index 332) (SNMP ifIndex 528)
  Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
  Input packets : 0
  Output packets: 13
  Protocol inet, MTU: 1500

```

### show interfaces af-interface-name brief

```

user@router> show interfaces af4 brief

Physical interface: af4, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1518, Clocking: Unspecified,
  Speed: 480000mbps
  Device flags   : Present Running
  Interface flags: Internal: 0x4000

  Logical interface af4.1
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ] Encapsulation: ENET2

```

```
inet 1.1.1.1/24
multiservice
```

### show interfaces af-interface-name extensive

```
user@router> show interfaces af4 extensive
```

```
Physical interface: af4, Enabled, Physical link is Up
  Interface index: 139, SNMP ifIndex: 514, Generation: 142
  Type: Ethernet, Link-level type: Ethernet, MTU: 1518, Clocking: Unspecified,
Speed: 480000Mbps
  Device flags   : Present Running
  Interface flags: Internal: 0x4000
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0,
state: unsuppressed
  Current address: 2c:6b:f5:55:eb:f6, Hardware address: 2c:6b:f5:55:eb:f6
  Alternate link address: Unspecified
  Last flapped   : 2017-10-18 19:40:00 EDT (02:50:04 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   :           4048           0 bps
    Output bytes  :          144092          59440 bps
    Input packets :             88           0 pps
    Output packets:            186           4 pps
  IPv6 transit statistics:
    Input bytes   :             0
    Output bytes  :             0
    Input packets :             0
    Output packets:             0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:
0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0
  Bandwidth      : 480 Gbps
  Peer GNF id    : 4
  Peer GNF Forwarding element(FE) view :
      FPC slot:FE Num      FE Bandwidth(Gbps)      Status
      6:0                  240                      Up
      6:1                  240                      Up

Logical interface af4.1 (Index 328) (SNMP ifIndex 593) (Generation 137)
  Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ] Encapsulation: ENET2
  Traffic statistics:
    Input bytes   :           414
    Output bytes  :          139906
    Input packets :             9
    Output packets:            107
  Local statistics:
    Input bytes   :           414
    Output bytes  :           598
    Input packets :             9
    Output packets:            13
  Transit statistics:
    Input bytes   :             0           0 bps
    Output bytes  :          139308          59240 bps
    Input packets :             0           0 pps
```

```

Output packets:          94          4 pps
Protocol inet, MTU: 1500
Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1, Curr new hold
cnt: 0, NH drop cnt: 0
Generation: 162, Route table: 0
Flags: Sendbroadcast-pkt-to-re
Output Filters: f-basic-sr-tcm-ca
Addresses, Flags: Is-Preferred Is-Primary
Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255, Generation:
148
Protocol multiservice, MTU: Unlimited, Generation: 163, Route table: 0
Policer: Input: __default_arp_policer__

```

### show interfaces af-interface-name detail

```
user@router> show interfaces af4 detail
```

```

Physical interface: af4, Enabled, Physical link is Up
Interface index: 139, SNMP ifIndex: 514, Generation: 142
Type: Ethernet, Link-level type: Ethernet, MTU: 1518, Clocking: Unspecified,
Speed: 480000mbps
Device flags   : Present Running
Interface flags: Internal: 0x4000
Link type      : Full-Duplex
Physical info   : Unspecified
Hold-times     : Up 0 ms, Down 0 ms
Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0,
state: unsuppressed
Current address: 2c:6b:f5:55:eb:f6, Hardware address: 2c:6b:f5:55:eb:f6
Alternate link address: Unspecified
Last flapped   : 2017-10-18 19:40:00 EDT (02:55:48 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes   :          77518216910          1944810944 bps
Output bytes  :          3059154          68832 bps
Input packets:          52450649          164494 pps
Output packets:          4672          20 pps
IPv6 transit statistics:
Input bytes   :          0
Output bytes  :          0
Input packets:          0
Output packets:          0
Bandwidth     : 480 Gbps
Peer GNF id   : 4
Peer GNF Forwarding element(FE) view :
      FPC slot:FE Num      FE Bandwidth(Gbps)      Status
      6:0                240                Up
      6:1                240                Up

Logical interface af4.1 (Index 328) (SNMP ifIndex 593) (Generation 137)
Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ] Encapsulation: ENET2
Traffic statistics:
Input bytes   :          77518213184
Output bytes  :          3054342
Input packets:          52450568
Output packets:          4591
Local statistics:
Input bytes   :          460
Output bytes  :          4600
Input packets:          10

```

```

Output packets:          100
Transit statistics:
Input bytes  :          77518212724          1944810944 bps
Output bytes :          3049742          68632 bps
Input packets:          52450558          164494 pps
Output packets:          4491          20 pps
Protocol inet, MTU: 1500
Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1, Curr new hold
cnt: 0, NH drop cnt: 0
Generation: 162, Route table: 0
Flags: Sendbroadcast-pkt-to-re
Output Filters: f-basic-sr-tcm-ca
Addresses, Flags: Is-Preferred Is-Primary
Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255, Generation:
148
Protocol multiservice, MTU: Unlimited, Generation: 163, Route table: 0
Policer: Input: __default_arp_policer__

```

### show interfaces af-interface-name statistics

This output is applicable to Junos OS Release 18.3R1 onwards.

```
user@router> show interfaces af4 statistics
```

```

Physical interface: af4, Enabled, Physical link is Up
Interface index: 165, SNMP ifIndex: 958
Type: Ethernet, Link-level type: Flexible-Ethernet, MTU: 1522, Speed: 240000mbps

Device flags : Present Running
Interface flags: Internal: 0x4000
Link type : Full-Duplex
Current address: 00:90:69:c7:31:62, Hardware address: 00:90:69:c7:31:62
Last flapped : 2018-08-07 21:47:10 PDT (00:58:48 ago)
Statistics last cleared: Never
Input rate : 1523826080 bps (2976222 pps)
Output rate : 112168 bps (232 pps)
Input errors: 0, Output errors: 0
Bandwidth : 240 Gbps
Peer GNF id : 4
Peer GNF Forwarding element(FE) view :
FPC slot:FE num FE Bandwidth(Gbps) Status Transmit Packets Transmit
Bytes
12:0 120 Up 2802
168120
12:1 120 Up 2848
170880

Residual Transmit Statistics :
Packets : 0 Bytes : 0

Fabric Queue Statistics :
FPC slot:FE num High priority(pkts) Low priority(pkts)
12:0 2802 0
12:1 2848 0
FPC slot:FE num High priority(bytes) Low priority(bytes)
12:0 168120 0
12:1 170880 0
Residual Queue Statistics :
High priority(pkts) Low priority(pkts)
0 0
High priority(bytes) Low priority(bytes)

```

0

0

```
Logical interface af4.1 (Index 335) (SNMP ifIndex 7354)
  Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ] Encapsulation: ENET2
  Input packets : 71572652
  Output packets: 5740
  Protocol inet, MTU: 1500
  Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1, Curr new hold
cnt: 0, NH drop cnt: 0
  Flags: Sendbcst-pkt-to-re
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 20.0.0/24, Local: 20.0.0.1, Broadcast: 20.0.0.255
  Protocol mpls, MTU: 1488, Maximum labels: 3
  Protocol multiservice, MTU: Unlimited

Logical interface af4.32767 (Index 336) (SNMP ifIndex 7355)
  Flags: Up SNMP-Traps 0x4004000 VLAN-Tag [ 0x0000.0 ] Encapsulation: ENET2
  Input packets : 0
  Output packets: 0
  Protocol multiservice, MTU: Unlimited
  Flags: None
```

## show system anomalies (GNF)

<b>Syntax</b>	<b>show system anomalies (all-anomalies   config   fru   system)</b>
<b>Release Information</b>	Command introduced in Junos OS Release 17.4R1.
<b>Description</b>	Display the incompatibilities between the software version running on the guest network function (GNF) and the version running on the base system (BSYS).
<b>Options</b>	<p><b>all-anomalies</b>—Display the multi-version software incompatibilities from all categories—system, configuration, and FRU.</p> <p><b>config</b>—Display the feature incompatibilities between software versions.</p> <p><b>fru</b>—Display the FRU-level incompatibilities between software versions. This can be an incompatibility pertaining to the support for a specific FRU.</p> <p><b>system</b>—Display the system-level incompatibilities between software versions. These include interprocess communication (IPC) message, CLI, or SNMP incompatibility.</p>
<b>Required Privilege Level</b>	View
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show system anomalies gnf-id (BSYS) on page 88</a></li> <li>• <a href="#">Generic Guidelines for Using JDM Server Commands on page 123</a></li> <li>• <a href="#">request virtual-network-functions on page 131</a></li> <li>• <a href="#">request server authenticate-peer-server on page 130</a></li> </ul>
<b>Output Fields</b>	<a href="#">Table 19 on page 112</a> lists the output fields for the <b>show system anomalies</b> command. Output fields are listed in the approximate order in which they appear.

*Table 19: show system anomalies gnf-id Output Fields*

Field Name	Field Description
<b>Anomaly Type</b>	Shows the software incompatibility type. The following are the possible values: <ul style="list-style-type: none"> <li>• SYS—Indicates system-specific incompatibilities.</li> <li>• FRU—Indicates FRU-specific incompatibilities.</li> <li>• CONFIG—Indicates feature-specific incompatibilities.</li> </ul>
<b>Anomaly ID</b>	Shows the incompatibility ID.



Table 19: show system anomalies gnf-id Output Fields (continued)

Field Name	Field Description
<b>Default Action</b>	Shows the default actions associated with incompatibilities. The following are the possible values: <ul style="list-style-type: none"> <li>WARN—Indicates the presence of a minor incompatibility. It causes a warning.</li> <li>ABORT—Indicates the presence of a major incompatibility. It causes an abort action.</li> </ul>
<b>Class</b>	Indicates of the incompatibility is present in GNF, or BSYS, or both.
<b>Message</b>	Incompatibility description.
<b>FRU ID</b>	Field-replaceable unit (FRU) ID. Applicable in the case of FRU-specific incompatibilities.
<b>Firmware</b>	Firmware of the FRU.

## Sample Output

```
user@router> show system anomalies
```

Anomaly Type	Anomaly ID FRU ID	Default Action Firmware	Class	Message
SYS incompatibility 1	100	WARN	GNF Present	system
SYS incompatibility 2	103	WARN	BOTH Present	system
SYS incompatibility 3	200	WARN	BSYS Present	system



## CHAPTER 10

# Configuration Statements for JDM

- [virtual-network-functions](#) on page 116
- [server](#) on page 117
- [interfaces \(Junos Node Slicing\)](#) on page 118
- [routing-options \(Junos Node Slicing\)](#) on page 119
- [system login \(Junos Node Slicing\)](#) on page 120
- [root-login \(JDM\)](#) on page 121

## virtual-network-functions

---

Syntax	<pre>virtual-network-functions <i>vnf-name</i> {     base-config <i>base-config</i> ;     chassis-type <i>chassis-type</i>;     id <i>id</i>;     resource-template <i>resource-template</i>; }</pre>
Hierarchy Level	[edit]
Release Information	Statement introduced in Junos OS Release 17.2R1.
Description	<p>Associate a GNF ID, base configuration, chassis type and resource template with the VNF.</p> <p>The GNFs that are configured and committed will appear as auto-complete options in operational commands.</p>
Options	<p><b><i>vnf-name</i></b>—Name of the VNF.</p> <p><b>base-config <i>base-config</i></b>—Choose a base configuration for the VNF.</p> <p><b>chassis-type <i>chassis-type</i></b>—Choose the type of the router chassis (for example, MX960) used as the base system (BSYS) in the node slicing setup.</p> <p><b>id <i>id</i></b>—Assign a GNF ID. <b>Range:</b> 1–10</p> <p><b>resource-template <i>resource-template</i></b>—Assign a resource template to the VNF. The server resource template defines the number of dedicated CPU cores and the size of DRAM to be assigned to a VNF.</p>
Required Privilege Level	root
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">server on page 117</a></li></ul>

## server

---

**Syntax**

```
server {
  interfaces {
    cb0 cb0-interface;
    cb1 cb1-interface;
    jdm-management jdm-management-interface;
    vnf-management gnf-management-interface;
  }
}
```

**Hierarchy Level** [edit]

**Release Information** Statement introduced in Junos OS Release 17.2R1.

**Description** Configure the server interfaces for the JDM and GNFs. These include a JDM management interface, a GNF management interface, and two server interfaces that are connected to the MX Series router.

**Options** **cb0 *cb0-interface***—The server interface that is connected to the control board 0 of the MX Series router.

**cb1 *cb1-interface***—The server interface that is connected to the control board 1 of the MX Series router.

**jdm-management *jdm-management-interface***—The server interface to be used as the JDM management interface.

**vnf-management *gnf-management-interface***—The server interface to be used as GNF management interface.

**Required Privilege Level** root

**Related Documentation**

- [virtual-network-functions on page 116](#)

## interfaces (Junos Node Slicing)

---

<b>Syntax</b>	<pre>interfaces <i>interface-name</i>{     unit <i>unit</i> {         family (inet inet6) {             address <i>address</i>;         }     } }</pre>
<b>Hierarchy Level</b>	[edit]
<b>Release Information</b>	Statement introduced in Junos OS Release 17.2.
<b>Description</b>	Configure an IP address for the JDM management interface.
<b>Options</b>	<p><b><i>interface-name</i></b>—Name of the interface.</p> <p><b><i>unit unit</i></b>—Interface unit number. This is a logical unit number. The only supported value is 0.</p> <p><b><i>family (inet inet6)</i></b>—Protocol family.</p> <ul style="list-style-type: none"><li>• <b>inet</b>—Indicates IPv4.</li><li>• <b>inet6</b>—Indicates IPv6.</li></ul> <p><b><i>address address</i></b>—IP address of the interface.</p>
<b>Required Privilege Level</b>	root
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">server on page 117</a></li><li>• <a href="#">virtual-network-functions on page 116</a></li></ul>

---

## routing-options (Junos Node Slicing)

---

<b>Syntax</b>	<pre>routing-options {   static {     route <i>route</i> {       next-hop <i>next-hop</i>;     }   } }</pre>
<b>Hierarchy Level</b>	[edit]
<b>Release Information</b>	Statement introduced in Junos OS Release 17.2.
<b>Description</b>	Configure static routes for JDM.
<b>Options</b>	<p><b>route <i>route</i></b>—Static route destination.</p> <p><b>next-hop <i>next-hop</i></b>—Specify the address of the next hop to destination.</p>
<b>Required Privilege Level</b>	root
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">server on page 117</a></li><li>• <a href="#">virtual-network-functions on page 116</a></li></ul>

## system login (Junos Node Slicing)

---

<b>Syntax</b>	<pre>system login {     user <i>username</i> ;     class <i>class-name</i>; }</pre>
<b>Hierarchy Level</b>	[root]
<b>Release Information</b>	Statement introduced in Junos OS Release 18.3 R1.
<b>Description</b>	Create a non-root user in JDM for Junos Node Slicing. The non-root user accounts function only inside JDM, not on the host server.
<b>Options</b>	<p><b>user <i>username</i></b>—Username of the account.</p> <p><b>class <i>class-name</i></b>—Predefined login classes that JDM supports for non-root users.</p> <ul style="list-style-type: none"><li>• super-user</li><li>• operator</li><li>• read-only</li><li>• unauthorized</li></ul>
<b>Required Privilege Level</b>	
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Non-Root Users in JDM (Junos Node Slicing) on page 61</a></li><li>• <a href="#">Configuring JDM on the x86 Servers on page 31</a></li></ul>



## root-login (JDM)

---

<b>Syntax</b>	<code>root-login (allow   deny );</code>
<b>Hierarchy Level</b>	[edit system services ssh]
<b>Release Information</b>	Statement introduced in Junos OS Release 18.3.
<b>Description</b>	Control user access to the JDM through SSH.
<b>Default</b>	<code>root-login allow</code> is the default.
<b>Options</b>	<p><b>allow</b>—Allow users to log in to the JDM as root through SSH.</p> <p><b>deny</b>—Disable users from logging in to the JDM as root through SSH. This configuration option is applicable only to the JDM management interface (jmgmt0). Setting this configuration option does not block the internal JDM to JDM communication, which uses root account with password-less authentication method.</p>
<b>Required Privilege Level</b>	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring JDM on the x86 Servers on page 31</a></li></ul>



## CHAPTER 11

# Operational Commands for JDM

- [Generic Guidelines for Using JDM Server Commands on page 123](#)
- [clear log \(JDM\)](#)
- [monitor list \(JDM\)](#)
- [monitor start \(JDM\)](#)
- [monitor stop \(JDM\)](#)
- [request server authenticate-peer-server](#)
- [request virtual-network-functions](#)
- [show virtual-network-functions](#)
- [show version vnf](#)
- [show version \(JDM\)](#)
- [show system cpu \(JDM\)](#)
- [show system storage \(JDM\)](#)
- [show system memory \(JDM\)](#)
- [show system network \(JDM\)](#)
- [restart \(JDM\)](#)

## Generic Guidelines for Using JDM Server Commands

---

The following are general guidelines on how to use the JDM server commands:

- Append **all-servers** to an operational command to take action on both the servers. Example: **request virtual-network-functions gnfl restart all-servers**.
- Append **server0** or **server1** to an operational command to take action on server0 or server1. Example: **request virtual-network-functions gnfl restart server0**.

By default, the operational commands work only on the local JDM.

- Use the **commit synchronize** command to ensure that the configuration committed on one server is synchronized with the other server. The synchronization is bidirectional. A JDM configuration change at either of the servers is synchronized with the other server. When a virtual machine (VM) is instantiated, the GNF-re0 VM instance starts on server0 and the GNF-re1 VM instance starts on server1.



NOTE: If you do not use the `commit synchronize` command, you must configure and manage the VMs on both the servers manually.

**Related  
Documentation**

- [show virtual-network-functions on page 132](#)
- [request virtual-network-functions on page 131](#)
- [request server authenticate-peer-server on page 130](#)

## clear log (JDM)

---

<b>Syntax</b>	<code>clear log <i>file-name</i></code>
<b>Release Information</b>	Command introduced in Junos OS Release 17.4R1.
<b>Description</b>	Clear the system log or trace files.
<b>Required Privilege Level</b>	View
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">monitor start (JDM) on page 127</a></li></ul>
<b>List of Sample Output</b>	<a href="#">clear log on page 125</a>
<b>Output Fields</b>	This command produces no output.

## Sample Output

### clear log

```
user@jdm> clear log syslog
```

## monitor list (JDM)

---

**Syntax**    `monitor list`

**Release Information**    Command introduced in Junos OS Release 17.4R1.

**Description**    Show status of monitored files.

**Required Privilege Level**    View

**Related Documentation**

- [monitor start \(JDM\) on page 127](#)

**List of Sample Output**    [monitor list on page 126](#)

### Sample Output

`monitor list`

```
user@jdm> monitor list
```

## monitor start (JDM)

<b>Syntax</b>	<b>monitor start</b> <i>file-name</i>
<b>Release Information</b>	Command introduced in Junos OS Release 17.4R1.
<b>Description</b>	Start displaying the system log or trace file and additional entries being added to those files.
<b>Additional Information</b>	Log files are generated by the routing protocol process or by system logging.
<b>Required Privilege Level</b>	View
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">monitor stop (JDM) on page 129</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">monitor start on page 127</a>
<b>Output Fields</b>	<a href="#">Table 20 on page 127</a> describes the output fields for the <b>monitor start</b> command. Output fields are listed in the approximate order in which they appear.

*Table 20: monitor start Output Fields*

Field Name	Field Description
<b>***filename***</b>	Name of the file from which entries are being displayed.
<b>Date and time</b>	Timestamp for the log entry.

## Sample Output

### monitor start

```

user@jdm> monitor start syslog

*** syslog ***
Oct 19 19:44:21 jdm mgd[4069]: UI_CMDLINE_READ_LINE: User 'root', command 'monit
    or start syslog '
Oct 19 19:44:29 jdm mgd[3268]: UI_CMDLINE_READ_LINE: User 'root', command 'delet
    e groups server0 server interfaces cb1 '
Oct 19 19:44:29 jdm mgd[3268]: UI_CFG_AUDIT_OTHER: User 'root' delete: [groups s
    erver0 server interfaces cb1]
Oct 19 19:44:35 jdm mgd[3268]: UI_CMDLINE_READ_LINE: User 'root', command 'delet
    e groups server1 server interfaces cb1 '
Oct 19 19:44:35 jdm mgd[3268]: UI_CFG_AUDIT_OTHER: User 'root' delete: [groups s

```

```
server1 server interfaces cb1]
Oct 19 19:44:36 jdm mgd[3268]: UI_CMDLINE_READ_LINE: User 'root', command 'commit '
Oct 19 19:44:36 jdm mgd[3268]: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Oct 19 19:44:36 jdm mgd[3268]: UI_COMMIT_PROGRESS: Commit operation in progress:
    Obtaining lock for commit
Oct 19 19:44:36 jdm mgd[3268]: UI_COMMIT_PROGRESS: Commit operation in progress:
    updating commit revision
Oct 19 19:44:36 jdm mgd[3268]: UI_COMMIT_PROGRESS: Commit operation in progress:
    obtaining db lock on server1
Oct 19 19:44:36 jdm mgd[3268]: UI_COMMIT_PROGRESS: Commit operation in progress:
    UI extensions feature is not configured
Oct 19 19:44:36 jdm mgd[3268]: UI_COMMIT_PROGRESS: Commit operation in progress:
    Started running translation script
```



## monitor stop (JDM)

---

**Syntax**    `monitor stop file-name`

**Release Information**    Command introduced in Junos OS Release 17.4R1.

**Description**    Stop displaying the system log or trace file.

**Additional Information**    Log files are generated by the routing protocol process or by system logging.

**Required Privilege Level**    View

**Related Documentation**    • [monitor start \(JDM\) on page 127](#)

**Output Fields**    This command produces no output.

## request server authenticate-peer-server

---

**Syntax**    `request server authenticate-peer-server`

**Release Information**    Command introduced in Junos OS Release 17.4R1.

**Description**    Copy the **ssh** public key to the peer JDM. This command is equivalent to **ssh-copy-id user@jdm-server<0/1>**.



**NOTE:** If the JDM fails to establish SSH connection with its peer on either of the two CB links, you need to run the JDM CLI command **request server authenticate-peer-server**. You can use the JDM CLI command **show server connections** to view the status of the SSH connection between the JDM peers. Note that the command **request server authenticate-peer-server** will prompt for user confirmation twice - once per CB link.

**Required Privilege Level**    View

**Related Documentation**

- [Generic Guidelines for Using JDM Server Commands on page 123](#)
- [show virtual-network-functions on page 132](#)
- [request virtual-network-functions on page 131](#)

## Sample Output

```
user@jdm> request server authenticate-peer-server
/usr/bin/ssh-copy-id: INFO: attempting to log in with the new key(s), to filter
out anythat are already installed
/usr/bin/ssh-copy-id: INFO: 1 key(s) remain to be installed -- if you are prompted
now it is to install the new keys
user@jdm-server1's password:
```

Number of key(s) added: 1

Now try logging into the machine, with:  "ssh 'user@jdm-server1'"  
and check to make sure that only the key(s) you wanted were added.

## request virtual-network-functions

**Syntax** `request virtual-network-functions vnf-name (add-image | all-servers | delete-image | force | restart | server | start | stop)`

**Release Information** Command introduced in Junos OS Release 17.2R1.  
The **force** option is available from Junos OS Release 17.4R1.

**Description** Start, stop or restart the VNFs. Also, you can add or remove the base image.



**NOTE:** You can issue these commands either on both the servers (`server0` and `server1`) or on one specific server.

**Options** *vnf-name*—Name of the VNF.

**add-image**—Add the base image to the VNF repository.

**all-servers**—Issue the command on both the servers.

**delete-image**—Delete the base image from the VNF repository.

**force**—Overwrite the image that was added earlier (example: `request virtual-network-functions vnf-name add-image new-image-name force`). The **force** option is also used with **delete-image** command to perform a cleanup in case, for example, you abruptly stopped an earlier **add-image** process by pressing Ctrl-C.

**restart**—Restart the VNF specified.

**server**—Issue the command on a specific server. Applicable value is **0** or **1**.

**start**—Start the VNF.

**stop**—Stop the VNF.

**Required Privilege Level** View

**Related Documentation**

- [Generic Guidelines for Using JDM Server Commands on page 123](#)
- [show virtual-network-functions on page 132](#)
- [request server authenticate-peer-server on page 130](#)

## show virtual-network-functions

**Syntax** `show virtual-network-functions (all-servers | server | vnf-name <detail>)`

**Release Information** Command introduced in Junos OS Release 17.2R1.

**Description** Display the list of guest network functions (GNFs) along with their IDs, status and availability.

**Options** **all-servers**—Display the details of the GNFs on both the servers.

**server**—Display the details of the GNFs on one specific server. Applicable value is **0** or **1**.

**vnf-name**—Display additional details of a particular GNF. You can use the **detail** option to view the detailed output. For example, `show virtual-network-functions gnf1 detail`.

**Required Privilege Level** View

**Related Documentation**

- [Generic Guidelines for Using JDM Server Commands on page 123](#)
- [request virtual-network-functions on page 131](#)
- [request server authenticate-peer-server on page 130](#)

**Output Fields** [Table 21 on page 132](#) lists the output fields for the `show virtual-network-functions` command.

*Table 21: show virtual network functions Output Fields*

Field Name	Field Description
ID	The ID associated with the VNF.
Name	Name of the VNF.
State	Status of the VNF. <ul style="list-style-type: none"> <li>• <b>Running</b>—The VNF is online and running.</li> <li>• <b>Shut off</b>—The VNF is in the shut down state.</li> </ul>
Liveness	Indicates the availability of VNF. <ul style="list-style-type: none"> <li>• <b>Up</b></li> <li>• <b>Down</b></li> </ul>
Name	Name of the VNF.
IP Address	IP address of the VNF.

Table 21: show virtual network functions Output Fields (continued)

Field Name	Field Description
Status	Status of a particular VNF. <ul style="list-style-type: none"> <li>• <b>Running</b>—The VNF is online and running.</li> <li>• <b>Shut off</b>—The VNF is in the shut down state.</li> </ul>
Cores	Number of cores in the VNF.
Memory	The VNF memory.
Resource Template	The resource template associated with a VNF.
Qemu Process id	Qemu process ID.
VNF CPU Utilization and Allocation Information	Shows the GNF CPU utilization details. See also: <a href="#">show system cpu (JDM)</a> .
VNF Memory Information	Displays the following memory information about the GNFs: <ul style="list-style-type: none"> <li>• Name—GNF name.</li> <li>• Resident—The memory used by the GNFs.</li> <li>• Actual—Actual memory.</li> </ul>
VNF Storage Information	Displays the following guest network function (GNF) storage information: <ul style="list-style-type: none"> <li>• Directories—Names of the directories.</li> <li>• Size—Total storage size.</li> <li>• Used—Storage used.</li> </ul>
VNF Interfaces Statistics	Shows the GNF interface statistics information. See also: <a href="#">show system network (JDM)</a> .
VNF Network Information	Shows the list of Physical Interfaces, Virtual Interfaces and MAC addresses.

## Sample Output

### show virtual-network-functions

```
user@jdm> show virtual-network-functions
```

ID	Name	State	Liveness
1	bittern-gnf-a	Running	Up
2	bittern-gnf-b	Running	Up
3	bittern-gnf-c	Running	Up
4	bittern-gnf-d	Running	Up
5	bittern-gnf-e	Running	Up

## Sample Output

### show virtual-network-functions (for a specific VNF)

```
user@jdm> show virtual-network-functions gnfl
```

#### Virtual Machine Information

```
-----
Name:          gnfl
IP Address:    192.168.2.1
Status:        Running
Liveness:      up
Cores:         4
Memory:        32GB
Resource Template: 4core-32g
Qemu Process id: 10891
```

## Sample Output

### show virtual-network-functions <vnf-name> detail

```
user@jdm> show virtual-network-functions gnfl detail
```

#### VNF Information

```
-----
ID          1
Name:       gnfl
Status:     Running
Liveness:   up
IP Address: 192.168.2.1
Cores:      2
Memory:     16GB
Resource Template: 2core-16g
Qemu Process id: 20478
SMBIOS version: v1
```

VNF Uptime: 206:09.58

#### VNF CPU Utilization and Allocation Information

GNF	CPU-Id(s)	Usage	Qemu Pid
gnfl	10,11	6.1%	20478

#### VNF Memory Information

Name	Actual	Resident
gnfl	16.0G	15.5G

#### VNF Storage Information

Directory	Size	Used
-----------	------	------

/vm-primary/gnf1

50.2G 4.2G

## VNF Interfaces Statistics

Interface		Rcvd Bytes		Rcvd packets		Rcvd Error		Rcvd Drop		Trxd	
bytes	Trxd	Packets	Trxd	Error	Trxd	Drop					
macvtap0				19077011		335687	0	0			
261601960	3545272	0	0								
macvtap1				99369778		689729	0	0			
172763580	2786344	0	0								
vnet1				24686		527	0	0		11620	
250	0	0									
macvtap2				428385		7405	0	0			
861020752	10813152	0	0								

## VNF Network Information

Virtual Interface	Physical Interface	MAC
macvtap0	ens3f0	00:85:49:54:cd:30
macvtap1	ens3f1	00:85:49:54:cd:31
vnet1	bridge_jdm_vm	00:85:49:54:cd:32
macvtap2	enp3s0f2	00:85:49:54:cd:33

## show version vnf

<b>Syntax</b>	<b>show version vnf <i>vnf-name</i></b>
<b>Release Information</b>	Command introduced in Junos OS Release 17.4R1.
<b>Description</b>	Display the hostname and version information about the specified guest network function (GNF).
<b>Options</b>	<b>vnf-name</b> —Name of the GNF for which you want to view the version details.
<b>Required Privilege Level</b>	View
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Generic Guidelines for Using JDM Server Commands on page 123</a></li> <li>• <a href="#">request virtual-network-functions on page 131</a></li> <li>• <a href="#">request server authenticate-peer-server on page 130</a></li> </ul>

## Sample Output

Depending on the platform running Junos OS, you might see different installed sub-packages.

### show version vnf

```
user@jdm> show version vnf gnf2

Hostname: gnf2
Model: mx960
Junos: 17.4X48-D10.3
JUNOS OS Kernel 64-bit [20170913.201739_fbsd-builder_stable_11]
JUNOS OS libs [20170913.201739_fbsd-builder_stable_11]
JUNOS OS runtime [20170913.201739_fbsd-builder_stable_11]
JUNOS OS time zone information [20170913.201739_fbsd-builder_stable_11]
JUNOS network stack and utilities [20170926.111120_builder_junos_174_x48_d10]
JUNOS modules [20170926.111120_builder_junos_174_x48_d10]
JUNOS mx modules [20170926.111120_builder_junos_174_x48_d10]
JUNOS libs [20170926.111120_builder_junos_174_x48_d10]
JUNOS OS libs compat32 [20170913.201739_fbsd-builder_stable_11]
JUNOS OS 32-bit compatibility [20170913.201739_fbsd-builder_stable_11]
JUNOS libs compat32 [20170926.111120_builder_junos_174_x48_d10]
JUNOS runtime [20170926.111120_builder_junos_174_x48_d10]
Junos vmguest package [20170926.111120_builder_junos_174_x48_d10]
JUNOS py extensions [20170926.111120_builder_junos_174_x48_d10]
JUNOS py base [20170926.111120_builder_junos_174_x48_d10]
JUNOS OS vmguest [20170913.201739_fbsd-builder_stable_11]
JUNOS OS crypto [20170913.201739_fbsd-builder_stable_11]
JUNOS mx libs compat32 [20170926.111120_builder_junos_174_x48_d10]
JUNOS mx runtime [20170926.111120_builder_junos_174_x48_d10]
JUNOS common platform support [20170926.111120_builder_junos_174_x48_d10]
JUNOS mx libs [20170926.111120_builder_junos_174_x48_d10]
```



JUNOS mtx Data Plane Crypto Support [20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS daemons [20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS mx daemons [20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS Services URL Filter package [20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS Services TLB Service PIC package [20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS Services SSL [20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS Services SOFTWARE [20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS Services Stateful Firewall [20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS Services RPM [20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS Services PTSP Container package [20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS Services PCEF package [20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS Services NAT [20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS Services Mobile Subscriber Service Container package  
[20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS Services MobileNext Software package  
[20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS Services Logging Report Framework package  
[20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS Services LL-PDF Container package [20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS Services Jflow Container package [20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS Services Deep Packet Inspection package  
[20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS Services IPSec [20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS Services IDS [20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS IDP Services [20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS Services HTTP Content Management package  
[20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS Services Crypto [20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS Services Captive Portal and Content Delivery Container package  
[20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS Services COS [20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS AppId Services [20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS Services Application Level Gateways  
[20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS Services ACL Container package [20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS SDN Software Suite [20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS Extension Toolkit [20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS Packet Forwarding Engine Support (wrlinux)  
[20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS Packet Forwarding Engine Support (MX/EX92XX Common)  
[20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS Packet Forwarding Engine Support (M/T Common)  
[20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS Packet Forwarding Engine Support (MX Common)  
[20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS jfirmware [20170926.111120\_builder\_junos\_174\_x48\_d10]  
JUNOS Online Documentation [20170926.111120\_builder\_junos\_174\_x48\_d10]

## show version (JDM)

---

<b>Syntax</b>	<b>show version</b> ( <b>all-servers</b>   <b>server</b>   <b>vnf</b>   <b>all-vnfs</b>   <b>detail</b>   <b>brief</b> )
<b>Release Information</b>	Command introduced in Junos OS Release 17.4R1.
<b>Description</b>	Display the version information about the Juniper Device Manager (JDM).
<b>Options</b>	<p><b>all-servers</b>—Display the version details of the JDM instances on both the servers.</p> <p><b>server</b>—Display the version details of the JDM instance on one specific server. <b>Range:</b> 0 through 1</p> <p><b>vnf</b>—Display the version details for a particular guest network function (GNF). You need to mention the GNF name in the command. Example: <b>show version vnf <i>gnf2</i></b>.</p> <p><b>all-vnfs</b>—Display the version details for all the GNFs.</p> <p><b>(detail   brief)</b>—Display the specified level of output.</p>
<b>Required Privilege Level</b>	View
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Generic Guidelines for Using JDM Server Commands on page 123</a></li><li>• <a href="#">request virtual-network-functions on page 131</a></li><li>• <a href="#">request server authenticate-peer-server on page 130</a></li></ul>

## Sample Output

### show version

```
user@jdm> show version

Hostname: mgb-dvaita-ixr1-jdm
Model: junos_node_slicing
Server slot : 1
JDM package version : 17.4-R1.7
Host Software [Red Hat Enterprise Linux]
JDM container Software [Ubuntu 14.04.1 LTS]
JDM daemon jdmd [Version: 17.4R1.7-secure]
JDM daemon jinventoryd [Version: 17.4R1.7-secure]
JDM daemon jdmmon [Version: 17.4R1.7-secure]
Host daemon jlinkmon [Version: 17.4R1.7-secure]

user@jdm> show version brief

Hostname: mgb-dvaita-ixr1-jdm
Model: junos_node_slicing
Server slot : 1
JDM package version : 17.4-R1.7
Host Software [Red Hat Enterprise Linux]
```

```
JDM container Software [Ubuntu 14.04.1 LTS]
JDM daemon jdmd [Version: 17.4R1.7-secure]
JDM daemon jinventoryd [Version: 17.4R1.7-secure]
JDM daemon jdmmon [Version: 17.4R1.7-secure]
Host daemon jlinkmon [Version: 17.4R1.7-secure]
```

```
user@jdm> show version detail
```

```
Hostname: mgb-dvaita-ixr1-jdm
Model: junos_node_slicing
Server slot : 1
JDM package version : 17.4-R1.7
Host Software [Red Hat Enterprise Linux]
JDM container Software [Ubuntu 14.04.1 LTS]
JDM daemon jdmd [Version: 17.4R1.7-secure]
JDM daemon jinventoryd [Version: 17.4R1.7-secure]
JDM daemon jdmmon [Version: 17.4R1.7-secure]
Host daemon jlinkmon [Version: 17.4R1.7-secure]
KERNEL 3.10.0-514.el7.x86_64
MGD release 17.4R1.7 built by builder on 2017-11-17 11:29:41 UTC
CLI release 17.4R1.7 built by builder on 2017-11-17 10:53:44 UTC
base-actions-dd release 17.4R1.7 built by builder on 2017-11-17 10:06:17 UTC
jdmd_common-actions-dd release 17.4R1.7 built by builder on 2017-11-17 10:06:09
UTC
jdmd_nv_jdm-actions-dd release 17.4R1.7 built by builder on 2017-11-17 10:06:09
UTC
```

## show system cpu (JDM)

<b>Syntax</b>	<code>show system cpu</code>
<b>Release Information</b>	Command introduced in Junos OS Release 17.4R1.
<b>Description</b>	Display the CPUs and their usage status.
<b>Required Privilege Level</b>	View
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Generic Guidelines for Using JDM Server Commands on page 123</a></li> <li>• <a href="#">request virtual-network-functions on page 131</a></li> <li>• <a href="#">request server authenticate-peer-server on page 130</a></li> </ul>
<b>Output Fields</b>	<a href="#">Table 22 on page 140</a> describes the output fields for the <b>show system cpu</b> command. Output fields are listed in the approximate order in which they appear.

*Table 22: show system cpu Output Fields*

Field Name	Field Description
VNF	Name of the GNF.
State	GNF status. For example, "Running".
CPU-Id(s)	List of CPU IDs.
Usage	CPU usage percentage.
Qemu Pid	Qemu process ID.

## Sample Output

### show system cpu

```

user@jdm> show system cpu

VNF CPU Utilization and Allocation Information
-----
VNF      -----          CPU-Id(s)          Usage  Qemu Pid
      State
-----
      -----

```

test 4,5,6,7,8,9,10,11 5.0% 32392

Running

Free CPUs : 12,13,14,15

Host Isolcpu(s): 2-15

Emulator Pins : 2-3

## show system storage (JDM)

<b>Syntax</b>	<code>show system storage</code>
<b>Release Information</b>	Command introduced in Junos OS Release 17.4R1.
<b>Description</b>	Display the JDM storage details such as storage size, used space, and available space.
<b>Required Privilege Level</b>	View
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Generic Guidelines for Using JDM Server Commands on page 123</a></li> <li>• <a href="#">request virtual-network-functions on page 131</a></li> <li>• <a href="#">request server authenticate-peer-server on page 130</a></li> </ul>
<b>Output Fields</b>	<a href="#">Table 23 on page 142</a> describes the output fields for the <b>show system storage</b> command. Output fields are listed in the approximate order in which they appear.

*Table 23: show system storage Output Fields*

Field Name	Field Description
Host Storage Information	Displays the following host storage information: <ul style="list-style-type: none"> <li>• Device—Host storage file system.</li> <li>• Size—Total storage size.</li> <li>• Used—Storage used.</li> <li>• Available—Storage available for use.</li> <li>• Use—Percentage of the storage being used.</li> <li>• Mount Point—Directory in which the host storage file system is mounted.</li> </ul>
JDM Storage Information	Displays the following JDM storage information: <ul style="list-style-type: none"> <li>• Directories—Names of the directories.</li> <li>• Used—Storage used.</li> </ul>
VNF Storage Information	Displays the following guest network function (GNF) storage information: <ul style="list-style-type: none"> <li>• Directories—Names of the directories.</li> <li>• Size—Total storage size.</li> <li>• Used—Storage used.</li> </ul>

## Sample Output

### show system storage

```
user@jdm> show system storage
```

## Host Storage Information

Device	Size	Used	Available	Use	Mount Point
/dev/mapper/rhel_mgb--dvaita--ixr0	493G	18G	451G	4%	/
/dev/sda2	976M	137M	772M	16%	/boot
/dev/sda1	200M	9.5M	191M	5%	/boot/efi
/dev/mapper/rhel_mgb--dvaita--ixr0	128G	40G	83G	33%	/var
/dev/mapper/rhel_mgb--dvaita--ixr0	99G	5.1G	89G	6%	/home

## JDM Storage Information

Directories	Used
/vm-primary/	8.0G
/var/third-party/	61M
/var/jdm-usr/	1.7G
/juniper	1.1G

## VNF Storage Information

Directories	Size	Used
/vm-primary/test	50.3G	4.5G
/vm-primary/test123	16.8G	3.5G

## show system memory (JDM)

<b>Syntax</b>	<code>show system memory</code>
<b>Release Information</b>	Command introduced in Junos OS Release 17.4R1.
<b>Description</b>	Display the memory usage information about the host server, Juniper Device Manager (JDM), and guest network functions (GNF).
<b>Required Privilege Level</b>	View
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Generic Guidelines for Using JDM Server Commands on page 123</a></li> <li>• <a href="#">request virtual-network-functions on page 131</a></li> <li>• <a href="#">request server authenticate-peer-server on page 130</a></li> </ul>
<b>Output Fields</b>	<a href="#">Table 24 on page 144</a> describes the output fields for the <b>show system memory</b> command. Output fields are listed in the approximate order in which they appear.

*Table 24: show system memory Output Fields*

Field Name	Field Description
Memory Usage Information	Displays the following memory usage information about host server and JDM: <ul style="list-style-type: none"> <li>• Total—Total memory.</li> <li>• Used—Used memory.</li> <li>• Free—Available memory.</li> </ul>
VNF Memory Information	Displays the following memory information about the GNFs: <ul style="list-style-type: none"> <li>• Name—GNF name.</li> <li>• Resident—The memory used by the GNFs.</li> <li>• Actual—Actual memory.</li> </ul>

## Sample Output

### show system memory

```

user@jdm> show system memory

Memory Usage Information
-----
      Total   Used   Free
-----
Host:  251G   14G   219G

JDM :   2.0G   2.0G    0

```



## VNF Memory Information

Name	Actual	Resident
gnf1	16.0G	13.6G
gnf2	16.0G	11.6G
gnf3	16.0G	12.6G
gnf4	16.0G	12.0G
gnf5	16.0G	12.7G

## show system network (JDM)

<b>Syntax</b>	<code>show system network</code>
<b>Release Information</b>	Command introduced in Junos OS Release 17.4R1.
<b>Description</b>	Display the statistics information for physical interface, JDM interface, and interfaces per guest network function (GNF).
<b>Required Privilege Level</b>	View
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Generic Guidelines for Using JDM Server Commands on page 123</a></li> <li>• <a href="#">request virtual-network-functions on page 131</a></li> <li>• <a href="#">request server authenticate-peer-server on page 130</a></li> </ul>
<b>Output Fields</b>	<a href="#">Table 25 on page 146</a> describes the output fields for the <code>show system network</code> command. Output fields are listed in the approximate order in which they appear.

*Table 25: show system network Output Fields*

Field Name	Field Description
<b>Physical Interfaces</b>	
Name	Name of the physical interface.
Index	Interface index.
MTU	The maximum transmission unit.
Hardware-address	The physical interface hardware address.
Rcvd Bytes	The number of bytes received.
Rcvd Packets	The number of packets received.
Rcvd Error	The packets received with error.
Rcvd Drop	The packets dropped during reception.
Trxd Bytes	The number of bytes transmitted.
Trxd Packets	The number of packets transmitted.
Trxd Error	The packets transmitted with error.

Table 25: show system network Output Fields (continued)

Field Name	Field Description
Trxd Drop	The packets dropped during transmission.
Flags	Interface flags.
<b>Per VNF Interface Statistics</b>	
Name	List of the interfaces for each GNF.
Source	Interface source.
MAC Address	The MAC address of the interface.
Rcvd Bytes	The number of bytes received.
Rcvd Packets	The number of packets received.
Rcvd Error	The packets received with error.
Rcvd Drop	The packets dropped during reception.
Trxd Bytes	The number of bytes transmitted.
Trxd Packets	The number of packets transmitted.
Trxd Error	The packets transmitted with error.
Trxd Drop	The packets dropped during transmission.
<b>JDM Interface Statistics</b>	
Name	Name of the JDM interface.
Index	Interface index.
MTU	The maximum transmission unit.
Hardware-address	The JDM interface hardware address.
Rcvd Bytes	The number of bytes received.
Rcvd Packets	The number of packets received.
Rcvd Error	The packets received with error.
Rcvd Drop	The packets dropped during reception.
Trxd Bytes	The number of bytes transmitted.

Table 25: show system network Output Fields (continued)

Field Name	Field Description
Trxd Packets	The number of packets transmitted.
Trxd Error	The packets transmitted with error.
Trxd Drop	The packets dropped during transmission.
Flags	Interface flags.
VNF MAC Address Pool	
Start MAC Address	MAC address start value.
Range	Available MAC address range.

## Sample Output

### show system network

```
user@jdm> show system network
```

#### Physical Interfaces

Name	Index	MTU	Hardware-address	Rcvd Bytes	Rcvd Packets	Rcvd Error
Rcvd Drop	Trxd Bytes	Trxd Packets	Trxd Error	Trxd Drop	Flags	
-----	-----	-----	-----	-----	-----	-----
enp3s0f1	4	1500	00:25:90:b5:75:51	8787662837	51975964	0
538926	40009223	407379	0	0	BMPRU	
ens3f1	7	1500	3c:fd:fe:08:87:02	1019880532	16723722	0
11243028	19265494115	31971968	0	0	BMPRU	
ens3f0	3	1500	3c:fd:fe:08:87:00	5951717054	81330473	0
11226877	139135292735	124708008	0	0	BMPRU	
enp3s0f2	5	1500	00:25:90:b5:75:52	3343179197	40806691	0
461955	3449064446	12191724	0	0	BMRU	

#### Per VNF Interface Statistics

Interface	Source	MAC Address	Rcvd Bytes	Rcvd
packets Rcvd Error Rcvd Drop Trxd bytes	Trxd Packets	Trxd Error	Trxd Drop	
-----	-----	-----	-----	-----
VNF name: test				
macvtap0	ens3f0	00:f1:60:3d:20:23	91526	1697
0	1254382	16782	0	0
macvtap1	ens3f1	00:f1:60:3d:20:24	561558	3727
0	803396	12958	0	0
vnet1	bridge_jdm_vm	00:f1:60:3d:20:25	2888	32
0	2282	25	0	0
macvtap2	enp3s0f2	00:f1:60:3d:20:26	152563	1080
0	4472700	55664	0	0

#### JDM Interface Statistics

Name	Index	MTU	Hardware-address	Rcvd Bytes	Rcvd Packets	Rcvd Error	Rcvd
Drop	Trxd Bytes	Trxd Packets	Trxd Error	Trxd Drop	Flags		
bme1	1433	1500	52:54:00:21:20:2e	502730	4506	0	0
	477328	2619	0	0	BMRU		
jmgmt0	1439	1500	00:f1:60:3d:20:22	4991675	66429	0	2862
	100548	891	0	0	BMRU		
bme2	1435	1500	52:54:00:88:b5:dd	2930	33	0	0
	3466	39	0	0	ABMRU		
cb0.4002	2	1500	00:f1:60:3d:20:20	12204921	209269	0	0
	3688591023	195579	0	0	ABMRU		
cb1.4002	3	1500	00:f1:60:3d:20:21	161850	3026	0	0
	204784	3029	0	0	ABMRU		

VNF MAC Address Pool

Start MAC Address: 00:f1:60:3d:20:20  
Range: 96

.....

## restart (JDM)

---

**Syntax** restart (gracefully | immediately | soft)  
<jdmd>  
<jdmmon>  
<jinventoryd>  
<jlinkmon>

**Release Information** Command introduced in Junos OS Release 17.4R1.

**Description** Restart daemons belonging to Juniper Device Manager (JDM).



**CAUTION:** Never restart a software process unless instructed to do so by a customer support engineer.

---

**Options** **gracefully**—(Optional) Restart a specified module or component by sending the equivalent of a UNIX SIGTERM signal.

**immediately**—(Optional) Immediately restart a module or component by sending the equivalent of a UNIX SIGKILL signal.

**soft**—(Optional) Reread and reactivate the configuration without completely restarting a module or component. This option is the equivalent of a UNIX SIGHUP signal.

**jdmd**—Restart the JDM service process.

**jdmmon**—Restart the JDM link monitoring daemon.

**jinventoryd**—Restart the JDM inventory management daemon.

**jlinkmon**—Restart the JDM link monitor daemon, which runs on the Linux host.



**NOTE:** The options **gracefully**, **immediately**, and **soft** are not available for restarting the Juniper link monitor daemon.

---

**Related Documentation**

- [Generic Guidelines for Using JDM Server Commands on page 123](#)
- [request virtual-network-functions on page 131](#)
- [request server authenticate-peer-server on page 130](#)

**List of Sample Output** [restart jdmd gracefully on page 151](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

restart jdmd gracefully

```
user@jdm> restart jdmd gracefully
Juniper Device Manager service process started, pid 12058
```

