



Juniper NETWORKS ACX7000 Series Cloud Metro Routers User Guide

[Home](#) » [JUNIPER NETWORKS](#) » Juniper NETWORKS ACX7000 Series Cloud Metro Routers User Guide 

Contents

- [1 Juniper NETWORKS ACX7000 Series Cloud Metro Routers](#)
- [2 Specifications](#)
- [3 Begin](#)
- [4 Up and Running](#)
- [5 Keep Going](#)
- [6 FAQ](#)
- [7 Documents / Resources](#)
 - [7.1 References](#)



Juniper NETWORKS ACX7000 Series Cloud Metro Routers



Specifications

- **Product Name:** Juniper Paragon Automation
- **Functionality:** End-to-end transport network automation
- **Supported Devices:** ACX7000 Series, PTX Series, MX Series, EX Series Switches, QFX Series Switches, and Cisco devices listed in Paragon Automation Supported Hardware

Begin

SUMMARY

This guide walks you through the simple steps to install Juniper® Paragon Automation and use Juniper® Paragon Automation to onboard, manage, and monitor networks devices.

Meet Paragon Automation

- Paragon Automation provides end-to-end transport network automation and simplifies the adoption of network automation for device, network, and service life cycles from Day 0 to Day 2.
- You can onboard ACX7000 Series, PTX Series, MX Series, EX Series Switches, QFX Series Switches, and Cisco devices listed in Paragon Automation Supported Hardware to Paragon Automation and manage them.

Install Paragon Automation

Before you install the Paragon Automation application, ensure that your server(s) meet the requirements listed in this section. A Paragon Automation cluster should contain only four nodes [virtual machines (VMs)], with three nodes acting as both primary and worker nodes and one node acting as a worker-only node.

Requirements

Hardware Requirements

The bare minimum resources required for each of the four nodes in the cluster are:

- 16-core vCPU
- 32-GB RAM
- 300-GB SSD (SSDs are mandatory)

NOTE:

- These VMs do not need to be in the same server, but the nodes need to be able to communicate over an L2 or L3 network.
- The hardware resources needed for each node VM depend on the size of the network that you want to onboard. To get a scale and size estimate of a production deployment and to discuss detailed dimensioning requirements, contact your Juniper Partner or Juniper Sales Representative.

Software Requirements

Use any of the following bare metal hypervisors to deploy Paragon Automation on one or more servers:

- VMware ESXi 8.0
- Kernel-based virtual machines (KVM)

For a KVM, use RHEL 8.10 as the host OS. For details about using KVM, see [Software Requirements](#).

- Proxmox Virtual Environment (VE)

Network Requirements

You can configure the Paragon Automation cluster by using IPv6 addresses in addition to IPv4 addresses. While configuring IPv6 addresses is optional, you must configure IPv4 addresses. The nodes can be in the same network or in the different networks. In either case, the four nodes of a Paragon Automation installation must be able to communicate with each other through SSH. You must have the following addresses available for the installation.

- IPv4 addresses (and optionally IPv6 addresses), one for each of the four nodes.
- A Virtual IP (VIP) address for generic ingress shared between gNMI, OC-TERM, and the Web UI. Alternatively, you can also use two VIP addresses—one for the Web GUI and another for gNMI and OC-TERM.
- A VIP address for Paragon Active Assurance Test Agent gateway (TAGW).
- A VIP address to establish Path Computational Element Protocol (PCEP) sessions between Paragon Automation and the devices for collecting label-switched path (LSP) information from the device.
- A VIP address (cRPD VIP) for establishing BGP Monitoring Protocol (BMP) session with external devices for routing observability.

NOTE:

- IPv6 address is not supported for the PCE server and establishing BMP sessions.
- You must configure the IPv6 addresses when you deploy the cluster. You cannot configure IPv6 addresses after a cluster is deployed using only IPv4 addresses.

If your cluster nodes are in different subnets, in addition to the listed addresses, you must ensure that:

- The cluster nodes have BGP connectivity with the respective upstream gateway top-of-rack (ToR).
- BGP peering is established between the cluster nodes and the ToR routers.

For more information, see [Network Requirements](#).


Browser Requirements

Paragon Automation is supported on the latest version of Google Chrome, Mozilla Firefox, and Safari.

Installation Workflow

- A system administrator can install Paragon Automation by downloading an OVA bundle and using the OVA bundle to deploy the node VMs on one or more VMware ESXi servers. Alternatively, you must extract the VMDK files from the OVA bundle and use them to deploy the node VMs on a KVM server or Proxmox VE. Paragon Automation runs on a Kubernetes cluster with three primary/worker nodes and one worker-only node. The installation is air-gapped but you need Internet access to download the OVA bundle to your computer.
- Figure on page 4 shows the workflow at a high-level for installing Paragon Automation.

Figure 1: Workflow for Installing Paragon Automation



For a detailed workflow, see [Paragon Automation Installation Workflow](#).

To install Paragon Automation, perform the following tasks:

1. Download the OVA bundle from the [Software Download Site](#).
2. **Create the VMs.**
 - You use the OVA (or OVF and VMDK files) bundle to create your node VMs. The software download files come prepackaged with the OS and all packages required to create the VMs and deploy your Paragon Automation cluster. The VMs have Ubuntu 22.04.5 LTS (Jammy Jellyfish) Linux base OS.
 - For details, see [Create the Node VMs](#).
3. **Configure the node VMs.**
 - Once the VMs are created, you must configure the hostname, IP address, DNS, and NTP server on each VM in the same way.
 - For details, see [Configure the Node VMs](#).
4. **Deploy the cluster.**

After all the VMs are configured, you can deploy the Paragon Automation cluster from the first VM. For details, see [Deploy the Cluster Nodes](#).

Log in to Paragon Automation

To log in to the Paragon Automation Web GUI:

1. Enter the common ingress VIP address in a browser to open the Paragon Automation login page.
 - The common ingress IP address, that you configured during installation, can be either IPv4 or IPv6.
 - To use the IPv4 address to connect to the Web GUI, enter the address in the https://ingress-vip format in

the URL. For example, `https://10.1.2.7`.

- To use the IPv6 address to connect to the Web GUI, enter the address in the `https://[ingress-vip-ipv6]` format in the URL. Ensure that you enclose the IPv6 address within square brackets. For example, `https://[2001:db8:1:2::7]`.
- Alternatively, if you have configured hostnames, you can use `https://ingress-vip-dns-hostname` to access the GUI.

2. Enter the Web admin user e-mail address and password that you configured while deploying Paragon Automation. The New Account page appears. You are now logged into Paragon Automation. You can now create organizations, sites, and users.

Add an Organization, a Site, and Users

Add an Organization

After you log in to the Paragon Automation GUI for the first time after installation, you must create an organization. After you create the organization, you are the superuser for the organization.

NOTE: You can add only one organization in this release. Adding more than one organization can lead to performance issues and constrain the disk space in the Paragon Automation cluster.

To create an organization:

1. Click Create Organization on the New Account page that appears after you log in to Paragon Automation. The Create Organization page appears.
2. Enter a name for the organization in Organization Name.
3. Click Create.

The organization is created. You are logged into the organization and the Troubleshoot Devices page appears. After you create an organization, you can add sites and users to the organization.

Create a Site

A site represents the location where devices are installed. You must be a superuser to add a site.

1. Click Inventory > Common Resources > Sites in the navigation menu.
2. On the Sites page, click + (Add) icon.
3. On the Create Site page, enter values for the fields Name, Location, Timezone, and Site Group.
4. Click Save.

The site is created and appears on the Sites page. For more information about sites, see [Add Sites](#).

Add Users

The superuser can add users and define roles for the users. To add a user to the organization:

1. On the banner, click Settings Menu > Users. The Users page appears.
2. Click the + (Invite User) icon. The New User page appears.
3. Enter the first name, surname, e-mail ID, and specify the role of the user in the Organization.
 - For the list of roles and their permissions in Paragon Automation, see [Predefined User Roles Overview](#).
 - The first name and surname can be up to 64 characters long.

4. Click Save.

- If SMTP is configured in Paragon Automation, an invite is sent to the user through an e-mail.
- If SMTP is not configured, the New User Creation page appears displaying the system-generated password for the user. You must share the password with the user manually.

5. (Optional) Follow step 1 through step 4 to add users with the Installer, Network Admin, and Observer roles.

Up and Running

SUMMARY

This section walks you through the preparatory steps that a Super User or Network Admin must perform before onboarding a device and moving the device to production.

Add Network Resource Pools

A network resource pool defines values for network resources, such as IPv4 loopback addresses, interface IP addresses, and so on, that are assigned to the devices in your network during device onboarding and for provisioning services (L2VPN, L3VPN, and L2 circuit).

You can create a network resource pool in Paragon Automation in one of the following ways:

- By configuring the resource pool in the Paragon Automation GUI.
- By uploading JSON files to Paragon Automation.
- By using REST APIs.

This section guides you through the steps to add network resource pools from the Paragon Automation GUI. For information about adding resource pools by using JSON files or REST APIs, see [Add Resource Pools](#).

To configure network resource pools in the Paragon Automation GUI:

1. Click **Orchestration > Services > Resource Instances** in the navigation menu. The Resource Instances page appears.
2. Click the + (Add) icon above the Resource Instances table. The Add New Resource Instance page appears.
3. In the Add New Resource Instance page:
 - Enter a name for the resource instance in the Instance Name field. For example, vpn-resource.
 - Enter the name of the customer for whom you are creating the resource instance in the Customer field. For example, for-abc-corp. The default name is network-operator.
 - Select the type of resource that you want to create from the Resource Design field.
For device onboarding, you must create L3-Addr, L2-Addr, and Routing resource pools. Start by selecting any one of the resource designs (for example, select L3-Addr to create layer 3 IP address pools).
4. Click Create.
 - The resource instance is created and the Modify Resource-Instance-Name page appears. The Modify Resource-
 - The instance-Name page lists an editor with the parameters that you can configure for the resource. For example, for the L3-Addr resource instance, configure the IPv4 prefixes and loopback addresses that can be assigned to the devices.
 - Alternatively, you can upload a JSON file populated with the resource values by using the Upload option

on the top-right corner of the Resource Editor.

- See [Configure Resource Pools](#) for more details.

5. Click Proceed.

The Compare Resource Definition page appears displaying the resources you have added.

6. Verify the resources you have added and then click Save and Commit.

Paragon Automation generates a service order to create the resources.

7. Repeat step 2 through step 6 to add the other two resource pools (for example, L2-Addr and Routing resources).

Add a Label

Labels can be used to identify devices of the same type or role and can be used as a reference in a device profile. For example, you can tag all provider edge devices with the label PE. Then, within a device profile, you can define that BGP sessions or MPLS LSPs should be established with any other device with the same label. When a provider edge device is onboarded using this profile, it gets tagged with label PE and automatically configured to peer with all the other devices also tagged with the label PE. At the same time, all these other devices also get configured to peer with this new device.

To add a label:

1. Navigate to Inventory > Devices > Device and Interface Profiles.
2. On the Devices and Interface Profiles page, click Add > Labels. The Create Labels page appears.
3. On the Create Labels page, enter the Plan Name (name for network implementation plan) and Label. For example, acx-onboarding-plan for the plan name and provider-edge-devices for label,
4. Click Save. The label is created and listed on the Device and Interface Profiles page.

Add a Device Profile

A device profile defines global configuration elements that are added to the device during onboarding. The configuration elements include hostname, IP address of the loopback, router ID, AS number, and protocols such as BGP and PCEP.

Before you add device profiles, ensure that you have

- Configured labels in Paragon Automation.
- Defined the resource pools. See “Add Network Resource Pools” on page 7.

To add a device profile:

1. Navigate to Inventory > Devices > Device and Interface Profiles.
2. In the Device and Interface Profiles page, click Add > Device Profile to create a device profile.
3. Enter the required information as explained in [Add a Device Profile](#).
4. Click Save. The device profile is created and appears on the Device and Interface Profiles page.

Add an Interface Profile

An interface profile defines interface-specific configuration elements that are added to the device during onboarding, including the interface’s IP address, whether the interface will be used for management or Internet connectivity, or whether the interface will be running OSPF, IS-IS, LDP, or RSVP protocols.

To add an interface profile:

1. Navigate to Inventory > Devices > Device and Interface Profiles.
2. In the Device and Interface Profiles page, click Add > Interface Profile to create an interface profile.
3. In the Create Interface Profile page, enter the required parameters as explained in [Add an Interface Profile](#).
NOTE: Enable the Internet Connected option for interfaces that connect with the Internet. Enabling this option allows Paragon Automation to initiate connectivity tests from the ports on which the interface profile is applied. We recommend that you enable this setting when you add the profile because you cannot enable or modify it later. For more information, see section [Device Connectivity Data and Test Results](#).
4. Click Save.

The interface profile is created and appears on the Device and Interface Profiles page.

Add a Network Implementation Plan

- To onboard a device, and enable health, connectivity, and compliance monitoring of the device after onboarding, you must create a network implementation plan that includes the device.
- Network implementation plans define which device and interface profiles should be applied to a device or a group of devices during onboarding. The profiles define which interfaces to configure, which protocols to enable, which IP addresses to assign, and so on.

To add a network implementation plan:

1. Navigate to Inventory > Device Onboarding > Network Implementation Plan.
2. On the Network Implementation Plan page, do one of the following:
 - Select the implementation plan that was created automatically after you created the device plan (the name of the plan will be the plan name you entered in the device profile), and then click Edit (pen) icon.
 - Click + (Add) to create a new network implementation plan. If you create a new plan instead, the device profiles that you created before are not available for selection within the implementation plan.

3. To create a new network implementation plan, enter a name for the plan and select a device profile and an interface profile.

If you want to set a default device profile and interface profile for the plan, select the names from the drop-down lists. If you are editing an automatically generated implementation plan, the default interface and device profiles are already populated.

4. Click Next to add devices to the plan.
5. In the Devices section click + (Add).
6. On the Add Device page, enter values for the hostname, IPv4 address, site, serial number, device vendor, and model, and select the device profile.

The serial number is used to map the device to this profile when it is added to the inventory (during adoption which is described later), and the onboarding process is started. The hostname, and IPv4 address that you enter here, along with all the other attributes included in the selected profiles are configured on the device during onboarding.

7. Click Next to go to the Physical Ports tab.

In the physical ports section:

1. Click + (Add) to enter the interfaces to be configured during onboarding.
2. Enter the interface name (include the unit number), a description for the interface, the IPv4 address, and

select the interface profile. You can also enter instructions for the installer to follow when physically installing the device and connecting the cables. Also, the pluggable field describes which type of optical transceiver is required.

3. Click OK to close the interface's configuration. Repeat this step for all the interfaces that will be part of the onboarding.
4. When you are finished entering all the interfaces, click Next to go to the Chassis tab.
5. In the Chassis tab, enter details about the power supply modules, fans, linecards, and optics.
6. Click Done when you are finished.
7. Repeat the steps 6 and 7 as needed to include all the devices, and its interfaces that you want to onboard under this implementation plan.
8. Click Next after you finish adding all the devices to the network implementation plan. The Links page appears.
9. Click + (Add) to add links between devices.
10. Click Next to view a summary of the configuration. If you want to modify the plan, you can click Edit and make the required changes.
11. Click Save.
 - The plan is created and appears on the Network Implementation Plan page.
 - For more information about adding a network implementation plan, see [Add a Network Implementation Plan](#).

Install a Device

A field technician should install the device at the site. For information about installing Juniper devices, see the Hardware guide of the respective device at <https://www.juniper.net/documentation/>. For installing Cisco Systems devices, refer to Cisco Systems documentation.

Onboard a Device

A superuser or network administrator can onboard a device by committing the outbound SSH commands to connect with Paragon Automation, on the device. This method of onboarding a device by committing the outbound SSH commands is also referred as "Adopting a Device".

You can onboard a device by any of the following methods:

- Onboard a device by using ZTP. In this method, you commit the SSH configuration on the device during ZTP.
- Onboard a device without ZTP. In this method, you manually commit the SSH configuration on the device.

For information about how to onboard a device, see the [Up and Running section](#) in the Onboard Juniper Networks Devices to Paragon Automation Quick Start Guide.

Approve a Device for Service

After a device is onboarded, a user with the superuser or network administrator can move the device to production and provision services on them.

To move a device to production:

1. Click Inventory > Device Onboarding > Onboarding Dashboard.
2. Filter the Ready for Service devices by selecting Ready for Service in the Operational State filter.
3. Click the Hostname link of the device to view the result of the automated tests that are performed on the Device-name page.

4. Analyze the results of the tests and view the alerts raised for the device. If no critical or major issues occur, you can move the device to production.
5. Click Put into Service to move the device to production. Paragon Automation changes the status of the device to In Service and moves the device to production. You can monitor the device for any alerts or alarms from the Device-Name (Observability > Troubleshoot Devices > Device-Name) page.

Keep Going

What's Next

Now that you've onboarded the device, here are some things you might want to do next.

If you want to	Then
Know how to troubleshoot alerts and alarms	See Troubleshoot Using Alerts and Alarms .
Find out how to provision and monitor a network service	See Service Orchestration
Know more about the device life cycle management use case	See Device Life Cycle Management Overview
Check trust and compliance of onboarded devices	See Perform Compliance Scans
Find out how to use active, synthetic traffic to monitor your network.	See Active Assurance

General Information

If you want to	Then
Use Paragon Automation to manage and monitor your devices.	See User Guide
Manage your Paragon Automation Account	See Manage your Paragon Automation Account
Learn about user roles in Paragon Automation	See Predefined User Roles Overview

If you want to	Then
Learn to manage, monitor, maintain, automate, and orchestrate network devices and services using Juniper Paragon Automation.	See Implementing Juniper Paragon Automation

Learn With Videos

If you want to	Then
Get short and concise tips and instructions that provide quick answers, clarity, and insight into specific features and functions of Juniper technologies.	See Learning with Juniper on Juniper Networks main YouTube page
View a list of the many free technical trainings we offer at Juniper.	Visit the Getting Started page on the Juniper Learning Portal.

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice. Copyright © 2025 Juniper Networks, Inc. All rights reserved.

FAQ

Q: How many nodes should be in a Paragon Automation cluster?

A: A Paragon Automation cluster should contain four nodes, with three nodes acting as primary and worker nodes, and one node as a worker-only node.

Q: Can I add multiple organizations in Paragon Automation?

A: No, you can add only one organization in this release to avoid performance issues and disk space constraints.

Documents / Resources

	Juniper NETWORKS ACX7000 Series Cloud Metro Routers [pdf] User Guide ACX7000 Series Cloud Metro Routers, ACX7000 Series, Cloud Metro Routers, Metro Routers, Routers
--	---

References

- [User Manual](#)

[Manuals+](#), [Privacy Policy](#)

This website is an independent publication and is neither affiliated with nor endorsed by any of the trademark owners. The "Bluetooth®" word mark and logos are registered trademarks owned by Bluetooth SIG, Inc. The "Wi-Fi®" word mark and logos are registered trademarks owned by the Wi-Fi Alliance. Any use of these marks on this website does not imply any affiliation with or endorsement.