

Quick Start

Juniper Routing Director 2.9.0 Quick Start

IN THIS GUIDE

- Step 1: Begin | 1
- Step 2: Up and Running | 8
- Step 3: Keep Going | 13

Step 1: Begin

SUMMARY

This quick start walks you through the simple steps to install Juniper® Routing Director (formerly Juniper® Paragon Automation) and use Juniper Routing Director to onboard, manage, and monitor network devices.

IN THIS SECTION

- Meet Routing Director | 1
- Install Routing Director | 2
- Log in to Routing Director | 6
- Add an Organization, a Site, and Users | 7

Meet Routing Director

Routing Director 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, Cisco devices, and Nokia 7750 SR-s devices listed in [Routing Director Supported Hardware](#) to Routing Director and manage them.

Install Routing Director

IN THIS SECTION

- [System Requirements | 2](#)
- [Install Routing Director on a Hypervisor | 4](#)
- [Install Routing Director on AWS | 5](#)

Before you install the Routing Director application, ensure that your server(s) meet the requirements listed in this section. You can deploy Routing Director as a multinode cluster of:

- Three nodes (all nodes acting as primary and worker nodes)
- Four nodes (three nodes acting as both primary and worker node)

For lab or demo environments, you can also deploy Routing Director on a single node that functions both as a primary and worker node.

You can deploy Routing Director on hypervisors and Amazon Web Service (AWS).

System Requirements

Hardware Requirements

You can deploy Routing Director in a single-node, three-node, or four-node cluster. [Table on page 2](#) lists the hardware requirements for each deployment.

Table 1: Hardware Requirements for Routing Director Deployment

Four-node cluster	Three-Node cluster	Single-node cluster
16-vCPU	24-vCPU	32-vCPU
32-GB RAM	48-GB RAM	64-GB RAM
512-GB SSD. SSDs are mandatory.	512-GB SSD. SSDs are mandatory.	512-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.
- The system requirements are different if you want to configure routing observability and AI/MLin Routing Director, see [Hardware Requirements](#) for details.

Software Requirements

You can deploy Routing Director on:

- One or more servers of the following bare-metal hypervisors:
 - VMware ESXi 8.0
 - Red Hat Enterprise Linux (RHEL) 8.10 and Ubuntu 22.04.05 kernel-based virtual machines (KVMs).
You must install the `libvirt`, `libvirt-daemon-kvm`, `bridge-utils`, and `qemu-kvm` packages. The hypervisor requires an x86-64 CPU,
 - Proxmox Virtual Environment (VE)
- AWS

See [Software Requirements](#).

Network Requirements

You can configure the Routing Director 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 nodes of a Routing Director 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 node.
- Virtual IP addresses for deployment on hypervisors:
 - 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 Active Assurance Test Agent gateway (TAGW).
 - A VIP address to establish Path Computational Element Protocol (PCEP) sessions between Routing Director 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.
- A VIP address (IPFIX VIP) for viewing details on predictor events (events indicating routing, forwarding, and OS exceptions that are identified as a potential indicator of traffic loss).

For deployment on AWS, you must input hostnames instead of IP addresses for generic ingress, TAGW, PCEP, cRPD VIP, and IPFIX.



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.
- IPv6 addresses are not supported in AWS deployment.

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


Routing Director is supported on the latest version of Google Chrome, Mozilla Firefox, and Safari.

Install Routing Director on a Hypervisor

A system administrator can install Routing Director 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. Routing Director 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 5](#) shows the workflow at a high-level for installing Routing Director.

Figure 1: Workflow for Installing Routing Director



For a detailed workflow, see [Install Routing Director](#).

To install Routing Director, 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 Routing Director 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 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 Routing Director cluster from the first VM. For details, see [Deploy the Cluster Nodes](#).

Install Routing Director on AWS

Deploy Routing Director on Amazon Web Services (AWS) by using a manual, nonintegrated model that runs the cluster on Amazon Elastic Compute Cloud (Amazon EC2) instances. This deployment model provides north-south access by using AWS Network Load Balancers (NLBs) in front of Kubernetes NodePort services.

Prerequisites to AWS Deployment

The following are the prerequisites to AWS deployment:

- Login credentials to AWS with permission to create, modify, and delete VMs, security groups, load balancers, IAM roles and policies, along with access to the S3 bucket.
- The S3 bucket must be preconfigured.

- The `vmimport` role must be created.
- A VPC subnet that can access the Internet on any NTP server.
- External DNS names for the following:
 - Hostname for generic ingress shared between the Routing Director Web GUI, NETCONF, and gNMI-term access.
 - Hostname for Paragon Active Assurance test agent gateway.
 - Hostname for PCE server.
 - Hostname for routing observability.



NOTE: Ensure that the administrator performing the installation is AWS Certified or has strong working knowledge of AWS.

To install and deploy Routing Director cluster on AWS:

1. Download Routing Director software OVA bundle from the software download site. Extract the VMDK files from the OVA bundle into Amazon Machine Image (AMI). See [Import Routing Director Image Files to AMI](#).
2. Create the security groups to allow traffic to and from the Routing Director cluster nodes. See [Create Security Groups](#).
3. Create and launch the EC2 VMs to provide the compute capacity and connectivity required for the Routing Director cluster. See [Provision AWS Compute and Networking Components](#).
4. Create the target groups and load balancers to enable access to the Routing Director cluster and its applications. See [Create the Target Groups and Load Balancers](#).
5. Deploy the cluster. See [Deploy the Cluster on AWS](#).

For detailed information, see [Install Juniper Routing Director on AWS](#).

Log in to Routing Director

To log in to the Routing Director Web GUI:

1. Enter the *common ingress VIP address* in a browser to open the Routing Director 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 Routing Director.

The New Account page appears. You are now logged into Routing Director. You can now create organizations, sites, and users.

Add an Organization, a Site, and Users

IN THIS SECTION

- [Add an Organization | 7](#)
- [Create a Site | 7](#)
- [Add Users | 8](#)

Add an Organization

After you log in to the Routing Director 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 Routing Director cluster.

To create an organization:

1. Click **Create Organization** on the New Account page that appears after you log in to Routing Director.
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**.
Alternatively, upload a CSV file containing site information for multiple sites to create them at once.
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 Routing Director, 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 Routing Director, 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.

Step 2: 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.

IN THIS SECTION

- [Add a Network Implementation Plan | 9](#)
- [Add Network Resource Pools | 11](#)
- [Add a Label | 12](#)
- [Install a Device | 12](#)
- [Onboard a Device | 12](#)
- [Approve a Device for Service | 13](#)

To onboard a brownfield device, follow the steps below:

1. Add a network Implementation plan. See ["Add a Network Implementation Plan" on page 9](#).
2. Install device. See ["Install a Device" on page 12](#).
3. Onboard the device. See ["Onboard a Device" on page 12](#).
4. Approve the device for service. See ["Approve a Device for Service" on page 13](#).

To onboard a greenfield device, follow the steps below:

1. Add a network Implementation plan. See ["Add a Network Implementation Plan" on page 9](#).
2. Add network resource pools. See ["Add Network Resource Pools" on page 11](#).
3. Add labels. See ["Add a Label" on page 12](#).
4. Install device. See ["Install a Device" on page 12](#).
5. Onboard the device. See ["Onboard a Device" on page 12](#).
6. Approve the device for service. See ["Approve a Device for Service" on page 13](#).

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 add the device to the plan. Network implementation plans define the configuration to be applied to a device or a group of devices during onboarding.

To add a network implementation plan:

1. Navigate to **Inventory > Device Onboarding > Network Implementation Plan**.
2. On the Network implementation Plan page, click + (Add) to create a network implementation plan.
The Create Network Implementation Plan page appears.
3. On the Create Network Implementation Plan page:
 - a. Enter a name for the plan.
 - b. Select one or more use cases that you want to configure on the device. The fields and tabs that appear on the Network Implementation Plan depend on the use cases you select here.
 - c. Add one or more default device profiles by clicking the **Add new device profile**. See [Add a Device Profile](#) for details

Configurations in the default device profiles are applied to all the devices added to the plan.
 - d. Add one or more default interface profiles by clicking the **Add new interface profile** See [Add an Interface Profile](#) for details.

Configurations in the default interface profiles are applied to all the interfaces added to the plan.
4. Click **Next** to add devices to the plan.
5. In the Devices section click + (Add).
6. On the General tab of the Add Device page, enter values for the hostname, site, serial number, device vendor, and model, and optionally create a device profile for the device. See [Add a Device Profile](#) for details.

The serial number is used to map the device to this profile when it is added to the inventory during the onboarding process. 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.

If you have selected Infrastructure and Service Orchestration use cases, configure the respective parameters.

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

In the Interfaces section:

- a. If the physical interfaces of the device are listed in the Interfaces table,
 - i. Select an interface and click the **Edit** icon to define the physical interface configuration on the Configure Port page.
 - ii. Select the interface and click **Set Managed** to configure the interface during onboarding.

If the device interfaces are not listed in the Interfaces table, click the **Add > Physical Interface** to add and configure the physical interface. The Configure Port page appears.

- b. On the Configure Port page, enter the interface name, select one or more device profiles, and port profiles..

If a port profile does not exist, click the **Add new port profile** link to create one or more port profiles to be applied to the physical interface during onboarding.

See [Add a Port Profile](#) for details.

- c. If Infrastructure Configuration is selected as a use case, enter a description, the name of the link the physical interface is connected to, port MTU, and the type of VLAN tagging for the interface in the Infrastructure Configuration section.

- d. If Service Orchestration is also selected as a use case, configure the access parameters.

- e. Click **OK** to close the Configure Port page.

- f. Repeat steps "7.a" on page 10 through "7.e" on page 10 for all the interfaces that you want to configure during device onboarding.

8. Add one or more logical interfaces to the physical interface.

- a. In the Interfaces table, select the physical interface for which you want to configure logical interfaces and click **Add > Logical Interface**.

- b. On the Configure Port page that appears, enter the logical unit number and add one or more interface profiles. If there is no suitable interface profile, create one by clicking the **Add new interface profile** link.

See [Add an interface Profile](#) for details.

- c. If Infrastructure Configuration is selected as a use case, enter a description for the logical interface, the IPv4 address with mask and IPv6 address with mask, and the VLAN ID.

- d. Click **OK** to return to the Interfaces table.

- e. Repeat steps "8.a" on page 10 through "8.d" on page 10 to add as many logical interfaces as you want.

9. Click **Done** to return to the Devices tab of the Add Network Implementation plan page.

10. Repeat the steps 5 through 9 to include all the devices that you want to onboard by using the implementation plan.

11. Click **Next** on the Add Network Implementation Plan page.

The Links tab appears

12. Click + (Add) to add links between devices added to the plan.

13. 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.

14. 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](#).

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 Routing Director in one of the following ways:

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

This section guides you through the steps to add network resource pools from the Routing Director 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 Routing Director 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-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**.

Routing Director 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 **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.

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 Routing Director, 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. After the configuration is committed, the device connects with Routing Director. The device is listed on the Inventory page with its status as Connected.

If the device is added to a network implementation plan, the configuration in the plan is committed on the device. You can view the status of device onboarding in the details of the network implementation plan (**Inventory > Device Onboarding > Network Implementation Plan > click a *Plan-Name* > More > Details**). After the device is onboarded, the onboarding status of the device is **Onboarded**.

For information about how to onboard a device, see [Add a Device to Routing Director](#).

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.

Routing Director 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.

Step 3: Keep Going

IN THIS SECTION

● [What's Next | 14](#)

- General Information | 14
- Learn With Videos | 15

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
Find out how to optimize your network.	See Network Optimization Overview
Find out more about planning and simulating network scenarios.	See Network Planner Overview

General Information

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

(Continued)

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

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.