Quick Start

Juniper Paragon Automation 2.2.0 Quick Start Guide

IN THIS GUIDE

Step 1: 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.

IN THIS SECTION

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 and Cisco Systems routers listed in Paragon Automation Supported Hardware to Paragon Automation and manage them.

Install Paragon Automation

IN THIS SECTION

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

Each node VM must have the following minimum hardware resources:

- 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 network.
- The hardware resources needed for each node VM depends 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 VMware ESXi 8.0 to deploy Paragon Automation.

Network Requirements

In this release, 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 four nodes of a Paragon Automation installation must be able to communicate with each other through SSH. You need to have the following addresses available for the installation, all in the *same IP network*.

- Four IP addresses, one for each of the four nodes
- Network gateway IP address
- A Virtual IP (VIP) address for generic ingress shared between gNMI, SSH ingress, and the Web UI.
- A VIP address for Paragon Active Assurance Test Agent gateway (TAGW).
- A VIP address (IPv4) to establish Path Computational Element Protocol (PCEP) sessions between Paragon Automation and the devices for collecting label-switched path (LSP) information from the device.

  NOTE: IPv6 address is not supported for the PCE server.

Browser Requirements

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

Create and Configure VMs

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 can also extract the OVF and VMDK files from the OVA bundle and use them to deploy the node VMs. 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 for installing Paragon Automation.

Figure 1: Workflow for Installing Paragon Automation

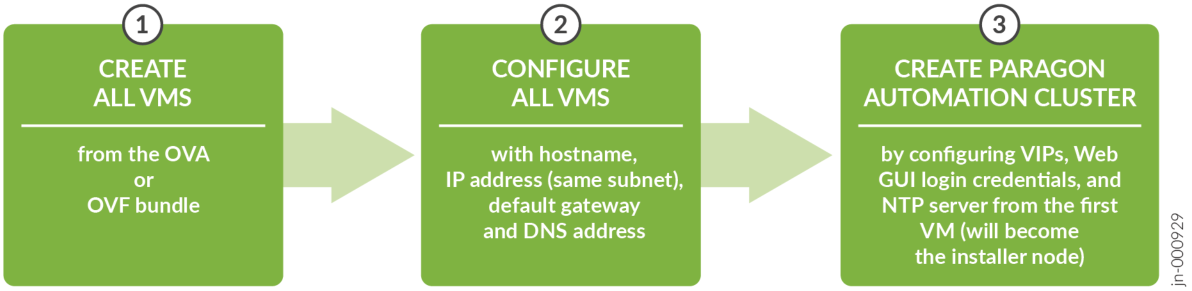
1
CREATE ALL VMS
from the OVA or OVF bundle
2
CONFIGURE ALL VMS
with hostname, IP address (same subnet), default gateway and DNS address
3
CREATE PARAGON AUTOMATION CLUSTER
by configuring VIPs, Web GUI login credentials, and NTP server from the first VM (will become the installer node)
jn-000929

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 a Linux base OS of Ubuntu 22.04.4 LTS (Jammy Jellyfish).

Once the VMs are created, you must configure each VM in the same way. When all the VMs are configured, you can deploy the Paragon Automation cluster from the first VM. Refer to Install Paragon Automation for a topology of the Paragon Automation cluster and the detailed Paragon Automation installation procedure.

To create a VM:

1. Download the OVA bundle onto your computer.

   You can use the OVA as a whole to create the VMs or alternatively, extract and use the OVF and .vmdk files from the OVA to create your VMs.

2. Log in to the VMware ESXi 8.0 server to install Paragon Automation.

3. Create the node VMs.

   To create the node VMs:

   a. Right-click the **Host** icon and select **Create/Register VM**.

      The New virtual machine wizard appears.

   b. On the Select creation type page, select **Deploy a virtual machine from an OVF or OVA file** and click **Next**.

   c. On the Select OVF and VMDK files page, enter a name for the node VM.

      Click to upload or drag and drop the OVA file or the OVF file along with the .vmdk files. Review the list of files to be uploaded and click **Next**.

   d. On the Select storage page, select the datastore that can accommodate 300-GB SSD for the node VM.

      Click **Next**. The extraction of files begins and takes a few minutes.

   e. On the Deployment options page:

      - Select the virtual network to which the node VM will be connected.

      - Select the Thick disk provisioning option.

- Enable the VM to power on automatically.

  Click **Next**.

f. On the Ready to complete page, review the VM settings and click **Finish** to create the node VM.

g. Power on the VM.

h. Follow steps "3.a" on page 4 to "3.g" on page 5 to create three more nodes. Enter appropriate VM names when prompted.

   Alternatively, if you are using VMware vCenter, you can right-click the VM, and click the **Clone > Clone to Virtual Machine** option to clone the newly created VM.

   Clone the VM thrice to create the remaining node VMs. Enter appropriate VM names when prompted.

i. After all the VMs are created, verify that the VMs have the correct specifications and are powered on.

4. Configure the Nodes.

   To configure the nodes:

   a. Connect to the node VM console of the first VM node.

      You are logged in to the node as the root user automatically and prompted to change your password.

   b. Enter and re-enter the new password.

      You are automatically logged out of the VM.

      > **NOTE**: We recommend that you enter the same password for all the VMs.

   c. When prompted, log in again as root user with the newly configured password.

   d. Configure the hostname and IP address of the VM, gateway, and DNS servers when prompted. You can also optionally choose to configure IPv6 addresses for your cluster.

      For information, see the *VM Configuration Wizard* table in Prepare the Nodes.

      > **NOTE**: For devices to be onboarded and managed by Paragon Automation, the devices must use either IPv4 addressing only or IPv6 addressing only to connect with Paragon Automation.

   e. When prompted, if you are sure to proceed, review the information displayed, type **y**, and press Enter.

   f. Repeat steps "4.c" on page 5 through "4.e" on page 5 for the other three VMs.

   g. After all the VMs are created, ensure that the nodes can reach each other.

      Type `exit` in a VM to go to the Linux root shell and ping each of the other three nodes.

You can now deploy the cluster.

Deploy the Cluster

Use the Paragon Shell CLI commands to deploy the Paragon Automation cluster.

To deploy a Paragon Automation cluster by using the Paragon Shell CLI commands:

1. Go back to the first node VM. If you have been logged out, log in again as the root user with the previously configured "password" on page 5.

   You are placed in Paragon Shell operational mode.

   ```
   *********************************************************************
                   WELCOME TO PARAGON SHELL!
        You will now be able to execute Paragon CLI commands!
   *********************************************************************
   root@eop>
   ```

2. Enter the configuration mode in Paragon Shell.

   ```
   root@eop> configure
   Entering configuration mode

   [edit]
   ```

3. Configure the following cluster parameters.

   ```
   root@eop# set paragon cluster nodes kubernetes 1 address 10.1.2.3

   [edit]
   root@eop# set paragon cluster nodes kubernetes 2 address 10.1.2.4

   [edit]
   root@eop# set paragon cluster nodes kubernetes 3 address 10.1.2.5

   [edit]
   root@eop# set paragon cluster nodes kubernetes 4 address 10.1.2.6

   [edit]
   root@eop# set paragon cluster ntp ntp-servers pool.ntp.org

   [edit]
   root@eop# set paragon cluster common-services ingress ingress-vip 10.1.2.7

   [edit]
   ```

```
root@eop# set paragon cluster applications active-assurance test-agent-gateway-vip 10.1.2.8

[edit]
root@eop# set paragon cluster applications web-ui web-admin-user "user-admin@juniper.net"

[edit]
root@eop# set paragon cluster applications web-ui web-admin-password Userpasswd

[edit]
```

Where:

- The IP addresses of `kubernetes` nodes with indexes 1 through 4 must match the static IP addresses configured on the node VMs. The Kubernetes nodes with indexes 1, 2, and 3 are the primary and worker nodes, the node with index 4 is the worker-only node.
- `ntp-servers` is the NTP server for synchronizing.
- `web-admin-user` and `web-admin-password` are the e-mail address and password that the first user can use to log in to the Web GUI.
- `ingress-vip` is the VIP address for the generic ingress IP address.
- `test-agent-gateway-vip` is the VIP address for the Paragon Active Assurance TAGW.

The VIP addresses are added to the outbound SSH configuration that is required for a device to establish a connection with Paragon Automation.

4. Configure the PCE server VIP address.

```
root@eop# set paragon cluster applications pathfinder pce-server pce-server-vip pce-server-vip
```

Where, *pce-server-vip* is the VIP address used by the PCE server to establish Path Computational Element Protocol (PCEP) sessions between Paragon Automation and the devices managed by it.

5. (Optional) If you had chosen to configure IPv6 addresses for your node VMs, enable IPv6 addressing and configure the IPv6 VIP addresses.

```
root@eop# set paragon cluster kubernetes address-family cluster-ipv6-enabled true

[edit]
root@eop# set paragon cluster common-services ingress ingress-vip-ipv6 2001:db8:1:2::7

[edit]
root@eop# set paragon cluster applications active-assurance test-agent-gateway-vip-ipv6 2001:db8:1:2::8

[edit]
root@eop# set paragon cluster install prefer-ipv6 true
```

```
[edit]
```

Where:

- `cluster-ipv6-enabled` enables usage of IPv6 addresses for the cluster making the cluster dual-stack.
- `ingress-vip-ipv6` is the IPv6 VIP address for generic common ingress and is used to connect to the Web GUI.
- `test-agent-gateway-vip-ipv6` is the IPv6 VIP address for the Active Assurance TAGW.
- `prefer-ipv6` configures preference for IPv6 addresses over IPv4 addresses. When set to `true`, and if hostnames are not configured, IPv6 VIP addresses are added to the outbound SSH configuration.

6. Configure hostnames for generic ingress and Paragon Active Assurance TAGW:

```
root@eop# set paragon cluster common-services ingress system-hostname ingress-vip-dns-hostname

[edit]
root@eop# set paragon cluster applications active-assurance test-agent-gateway-hostname nginx-ingress-
controller-hostname

[edit]
```

Where:

- `system-hostname` is the hostname for the generic ingress VIP address.
- `test-agent-gateway-hostname` is the hostname for the Paragon Active Assurance TAGW VIP address.

When you configure hostnames, the hostnames are added to the outbound SSH configuration instead of the VIP addresses. The hostnames can resolve to either IPv4 or IPv6 VIP addresses or both.

7. (Optional) Configure the following settings for SMTP-based user management.

```
root@eop# set paragon cluster mail-server smtp-relayhost smtp.relayhost.com

[edit]
root@eop# set paragon cluster mail-server smtp-relayhost-username relayuser

[edit]

root@eop# set paragon cluster mail-server smtp-relayhost-password relaypassword
[edit]

root@eop# set paragon cluster mail-server smtp-allowed-sender-domains paragonautomation.net

[edit]
```

```
root@eop# set paragon cluster mail-server smtp-sender-email no-reply@paragonautomation.net
[edit]
root@eop# set paragon cluster mail-server smtp-sender-name Juniper Paragon Automation

[edit]
root@eop# set paragon cluster papi papi-local-user-management false

[edit]
root@eop# set paragon cluster mail-server smtp-enabled true

[edit]
```

Where:

- *sender-domains* are the e-mail domains from which Paragon Automation sends e-mails to users.
- *relayhost-hostname* is the name of the SMTP server that relays messages.
- *relayhost-username* (optional) is the username to access the SMTP (relay) server.
- *relayhost-password* (optional) is the password for the SMTP (relay) server.
- *sender-e-mail-address* is the e-mail address that appears as the sender's e-mail address to the e-mail recipient.
- *sender-name* is the name that appears as the sender’s name in the e-mails sent to users from Paragon Automation.
- `papi-local-user-management false` disables local authentication.

NOTE:

- SMTP configuration is optional at this point. SMTP settings can be configured after the cluster has been deployed also. For information about how to configure SMTP after cluster deployment, see Configure SMTP Settings in Paragon Shell.
- For details about the behavior of Paragon Automation with different combinations of local authentication and SMTP configuration, see User Activation and Login.

8. (Optional) Install custom user certificates.

NOTE: Before you install user certificates, you must copy the custom certificate file and certificate key file to the **/root/epic/config** folder in the Linux root shell of the node from which you are deploying the cluster.

```
root@eop# set paragon cluster common-services ingress user-certificate use-user-certificate true

[edit]
```

```
root@eop# set paragon cluster common-services ingress user-certificate user-certificate-filename
"certificate.cert.pem"

[edit]
root@eop# set paragon cluster common-services ingress user-certificate user-certificate-key-filename
"certificate.key.pem"

[edit]
```

Where:

- *certificate.cert.pem* is the user certificate file name.
- *certificate.key.pem* is the user certificate key file name.

> **NOTE**: Installing certificates is optional at this point. You can configure Paragon Automation to use custom user certificates after cluster deployment also. For information about how to install user certificates after cluster deployment, see Install User Certificates.

9. Commit the configuration and exit configuration mode.

```
root@eop# commit
commit complete

[edit]
root@eop# exit
Exiting configuration mode

root@eop>
```

10. Generate the configuration files.

```
root@eop> request paragon config
Paragon inventory file saved at /epic/config/inventory
Paragon config file saved at /epic/config/config.yml
```

The **inventory** file contains the IP addresses of the VMs. The **config.yml** file contains minimum Paragon Automation cluster parameters required to deploy a cluster.

The `request paragon config` command also generates a **config.cmgd** file in the **config** directory. The **config.cmgd** file contains all the `set` commands that you executed in step "3" on page 6. If the **config.yml** file is inadvertently edited or corrupted, you can redeploy your cluster using the `load set config/config.cmgd` command in the configuration mode.

11. Generate SSH keys on the cluster nodes.

    When prompted, enter the SSH password for the VMs. Enter the same "password" on page 5 that you configured to log in to the VMs.

```
root@eop> request paragon ssh-key
Setting up public key authentication for ['10.1.2.3','10.1.2.4','10.1.2.5','10.1.2.6']

Please enter SSH username for the node(s): root
Please enter SSH password for the node(s):
                        password

checking server reachability and ssh connectivity ...
Connectivity ok for 10.1.2.3
Connectivity ok for 10.1.2.4
Connectivity ok for 10.1.2.5
Connectivity ok for 10.1.2.6
SSH key pair generated in 10.1.2.3
SSH key pair generated in 10.1.2.4
SSH key pair generated in 10.1.2.5
SSH key pair generated in 10.1.2.6
copied from 10.1.2.3 to 10.1.2.3
copied from 10.1.2.3 to 10.1.2.4
copied from 10.1.2.3 to 10.1.2.5
copied from 10.1.2.3 to 10.1.2.6
copied from 10.1.2.4 to 10.1.2.3
copied from 10.1.2.4 to 10.1.2.4
copied from 10.1.2.4 to 10.1.2.5
copied from 10.1.2.4 to 10.1.2.6
copied from 10.1.2.5 to 10.1.2.3
copied from 10.1.2.5 to 10.1.2.4
copied from 10.1.2.5 to 10.1.2.5
copied from 10.1.2.5 to 10.1.2.6
copied from 10.1.2.6 to 10.1.2.3
copied from 10.1.2.6 to 10.1.2.4
copied from 10.1.2.6 to 10.1.2.5
copied from 10.1.2.6 to 10.1.2.6
```

NOTE: If you have configured different passwords for the VMs, ensure that you enter corresponding passwords when prompted.

12. Deploy the cluster.

```
root@eop> request paragon deploy cluster
Process running with PID: 231xx03
To track progress, run 'monitor start /epic/config/log'
After successful deployment, please exit Paragon-shell and then re-login to the host to finalize the setup
```

The cluster deployment begins and takes over an hour to complete.

13. (Optional) Monitor the progress of the deployment onscreen.

```
root@eop> monitor start /epic/config/log
```

The progress of the deployment is displayed. Deployment is complete when you see an output similar to this onscreen.

```
<output snipped>

PLAY RECAP ********************************************************************************************
10.1.2.3               : ok=109  changed=33   unreachable=0    failed=0    rescued=0    ignored=0
10.1.2.4               : ok=34   changed=1    unreachable=0    failed=0    rescued=0    ignored=0
10.1.2.5               : ok=34   changed=1    unreachable=0    failed=0    rescued=0    ignored=0
10.1.2.6               : ok=30   changed=0    unreachable=0    failed=0    rescued=0    ignored=0

Monday 15 July 2024  18:56:14 +0000 (0:00:00.819)       0:01:23.328 ***********
===============================================================================
Unpack 3rdparty OS packages -------------------------------------------------------------------- 8.59s
Gathering Facts -------------------------------------------------------------------------------- 4.18s
add etcd user on running nodes ----------------------------------------------------------------- 3.85s
Gathering Facts -------------------------------------------------------------------------------- 3.61s
Gathering Facts -------------------------------------------------------------------------------- 3.04s
Gathering Facts -------------------------------------------------------------------------------- 2.98s
readonly : Synchronize k8s readonly config to rest of nodes ------------------------------------ 2.07s
kubernetes/addons/traefik : Install Helm Chart ------------------------------------------------- 1.82s
kubernetes/addons/traefik-paa : Install Helm Chart --------------------------------------------- 1.74s
Record installation status --------------------------------------------------------------------- 1.64s
kubernetes/addons/metadata : Create common/metadata -------------------------------------------- 1.38s
kubernetes/addons/traefik-paa : upload ingress spec files -------------------------------------- 1.30s
kubernetes/addons/traefik : upload ingress spec files ------------------------------------------ 1.28s
kubernetes/addons/rook-quota : get rook total capacity ----------------------------------------- 1.25s
kubernetes/addons/traefik : apply traefik ingress routes --------------------------------------- 1.09s
Record installation status --------------------------------------------------------------------- 0.98s
kubernetes/addons/traefik : apply traefik additional services ---------------------------------- 0.90s
kubernetes/addons/traefik : create default ingress cert ---------------------------------------- 0.89s
kubernetes/addons/metadata : Create config_yml ------------------------------------------------- 0.87s
```

```
kubernetes/addons/metadata : Get docker labels ------------------------------------------------- 0.86s
Playbook run took 0 days, 1 hours, 1 minutes, 23 seconds
registry-5749
root@eop>
```

Alternatively, if you did not choose to monitor the progress of the deployment onscreen using the `monitor` command, you can view the contents of the log file using the `file show /epic/config/log` command. The last few lines of the log file must look similar to "Sample Output" on page 12. We recommend that you check the log file periodically to monitor the progress of the deployment.

The CLI command prompt displays your login username and the node hostname that you configured previously. For example, if you entered Primary1 as the hostname of your primary node, the command prompt is `root@Primary1>`.

Upon successful completion of the deployment, the Paragon Automation cluster is created.

14. Log out of the VM and log in again to the Paragon Shell.

    The console output displays the Paragon Shell welcome message and the IP addresses of the four nodes (called Controller-1 through Controller-4), the Paragon Active Assurance TAGW VIP address, the Web admin user e-mail address, and Web GUI IP address. If you had enabled IPv6 addressing, the console output also displays the configured IPv6 addresses.

```
Welcome to Juniper Paragon Automation OVA



This Controller IP: 10.1.2.3, 2001:db8:1:2::3
This VM 10.1.2.3, 2001:db8:1:2::3 is part of an EPIC on-prem system with IPv6 enabled.
=======================================================================================
Controller IP    :  10.1.2.3, 10.1.2.4, 10.1.2.5, 10.1.2.6
PAA Virtual IP   :  10.1.2.8, 2001:db8:1:2::8
UI               :  https://10.1.2.7, https://[2001:db8:1:2::7]
Web Admin User   :  admin-user@juniper.net
=======================================================================================
ova: 20240503_2010
build: eop-release-2.2.0.6928.g6be8b6ce52

***************************************************************
                WELCOME TO PARAGON SHELL!
     You will now be able to execute Paragon CLI commands!
***************************************************************
root@Primary1>
```

You can now log in to the Paragon Automation GUI by using the Web admin user ID and password.

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

IN THIS SECTION

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 Super User 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 Super User 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 Super User 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 upto 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 Steps 1 through 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 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 UI. 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 Automaiton 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 design (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**.

   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 **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 16.

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:

   a. Click **+** (Add) to enter the interfaces to be configured during onboarding.

   b. 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.

   c. Click **OK** to close the interface's configuration. Repeat this step for all the interfaces that will be part of the onboarding.

   d. When you are finished entering all the interfaces, click **Next** to go to the Chassis tab.

   e. In the Chassis tab, enter details about the power supply modules, fans, linecards, and optics.

   f. Click **Done** when you are finished.

   g. 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 on 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 there are no critical or major issues, 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.

Step 3: Keep Going

IN THIS SECTION

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.
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 provision and monitor a network service	See Service Orchestration

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

(Continued)

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.