

# Juniper Apstra 5.0.1 / 5.0.0 Installation and Upgrade Guide

Published  
2025-03-10

# Table of Contents

## Apstra Installation

### Installation Requirements | 1

Installation Overview | 1

Supported Hypervisors and Versions | 2

Required Server Resources | 2

Required Communication Ports | 8

Additional Network Protocols | 9

Network Client Services | 9

### Install Apstra on ESXi | 10

### Install Apstra on KVM | 14

Install on KVM with Virtual Machine Manager | 14

Install on KVM with CLI | 18

### Install Apstra on Hyper-V | 20

### Install Apstra on VirtualBox | 23

### Configure Apstra Server | 25

Configure Apstra Server Static Management IP Address | 28

Change Apstra Server Hostname | 28

Configure New SSH Keys on Apstra Server | 28

Configure Apstra Docker Subnets | 29

Replace Apstra Server SSL Certificate | 30

Replace SSL Certificate on Apstra Server with Signed One | 31

Replace SSL Certificate on Apstra Server with Self-Signed One | 33

## Apstra Upgrade

### Supported Apstra Server Upgrade Paths | 35

### Apstra Server Upgrade Workflow | 36

## Upgrade Apstra on New VM (VM-VM) | 38

Step 1: Pre-Upgrade Validation | 38

Step 2: Deploy New Apstra Server | 40

Step 3: Import State | 41

Step 4: Import State for Intent-Based Analytics | 45

Step 5: Keep Old VM's IP Address (Optional) | 45

Step 6: Change Operation Mode to Normal | 46

Step 7: Upgrade Onbox Agents | 48

Step 8: Shut Down Old Apstra Server | 49

## Roll Back Apstra Server Upgrade | 50

# Apstra Installation

## IN THIS SECTION

- Installation Requirements | [1](#)
- Install Apstra on ESXi | [10](#)
- Install Apstra on KVM | [14](#)
- Install Apstra on Hyper-V | [20](#)
- Install Apstra on VirtualBox | [23](#)
- Configure Apstra Server | [25](#)

## Installation Requirements

### IN THIS SECTION

- Installation Overview | [1](#)
- Supported Hypervisors and Versions | [2](#)
- Required Server Resources | [2](#)
- Required Communication Ports | [8](#)
- Additional Network Protocols | [9](#)
- Network Client Services | [9](#)

## Installation Overview

Before installing Juniper Apstra software, refer to the following sections and ensure that the server where you'll install it meets requirements. Then you can install and configure Apstra on one of the supported hypervisors. Default passwords are not secure, so make sure to replace them with secure ones during configuration. We know how important complex passwords are, so we've made it a requirement to change default passwords with more complex ones than previous versions. We also recommend replacing the self-signed SSL certificate with a signed one from your own certificate authority so your environment is more secure. Keep reading for installation and configuration steps.

## Supported Hypervisors and Versions

Table 1: Supported Hypervisors

Hypervisor	Supported Versions
VMware ESXi	8.0, 7.0, 6.7, 6.5, 6.0
QEMU / KVM for Ubuntu	22.04 LTS, 18.04 LTS
Microsoft Hyper-V	Windows Server 2019 Datacenter Edition Windows Server 2016 Datacenter Edition
Oracle VirtualBox / VMware Workstation	For lab / evaluation purposes only

## Required Server Resources

### IN THIS SECTION

- [Resource Requirements | 2](#)
- [Apstra Cluster Node Examples | 3](#)

### Resource Requirements

Apstra server VM resource requirements are based on several factors:

- Size of the network
- Number of offbox agents
- Intent-Based Analytics (IBA) usage

Each VM node (controller, worker) supports up to 25 offbox agents. Make sure you allocate enough resources for your network. If one VM is insufficient for your needs, you can increase capacity by adding worker nodes with [Apstra VM Clusters](#).



**CAUTION:** If you allocate insufficient CPU and RAM allocations for your network you'll probably experience issues such as slow response times, system errors and a critical "segmentation fault" (core dump). If this happens, you'll need to delete the VM and redeploy it with additional resources.

The recommended resources are as follows:

**Table 2: Recommended VM Resources**

Resource	Recommendation
Memory	64 GB RAM + 500 MB per installed offbox agent*
CPU	8 vCPU
Disk	160 GB**
Network	1 network adapter, initially configured with DHCP

\* Container memory usage is dependent on the number of IBA collectors enabled. At a minimum, you'll need to [change the application weight](#) for Juniper offbox agents after installation is complete and you're in the Apstra environment.

\*\* Apstra images ship with an 80 GB disk by default. ESXi images ship with a second "empty" disk. On first boot, Apstra automatically runs aos\_extend\_disk, and if space is available, it extends / (root), /var, /var/log, and /var/log-aos/db to the new disk. (Shipping with an 80 GB disk instead of 160 GB keeps the image size reasonable.)

If you deploy Linux KVM QCOW2 or Microsoft Hyper-V VHDX, the second disk isn't included so the default is 80 GB. You can manually add an additional disk. Run aos\_extend\_disk yourself to extend / (root), /var, /var/log, and /var/log-aos/db to the new disk. For more information, see [Juniper Support Knowledge Base article KB37699](#).

### Apstra Cluster Node Examples

When you've installed your VMs and added offbox agents and IBA probes, you can monitor usage from the Apstra GUI. From the left navigation menu, navigate to **Platform > Apstra Cluster** to go to the Apstra Cluster Nodes Summary Table.

Our example includes one controller and two workers. The controller can usually accommodate any IBA requirement, but if you're using offbox agents you'll need an additional VM for every 25 or fewer agents.

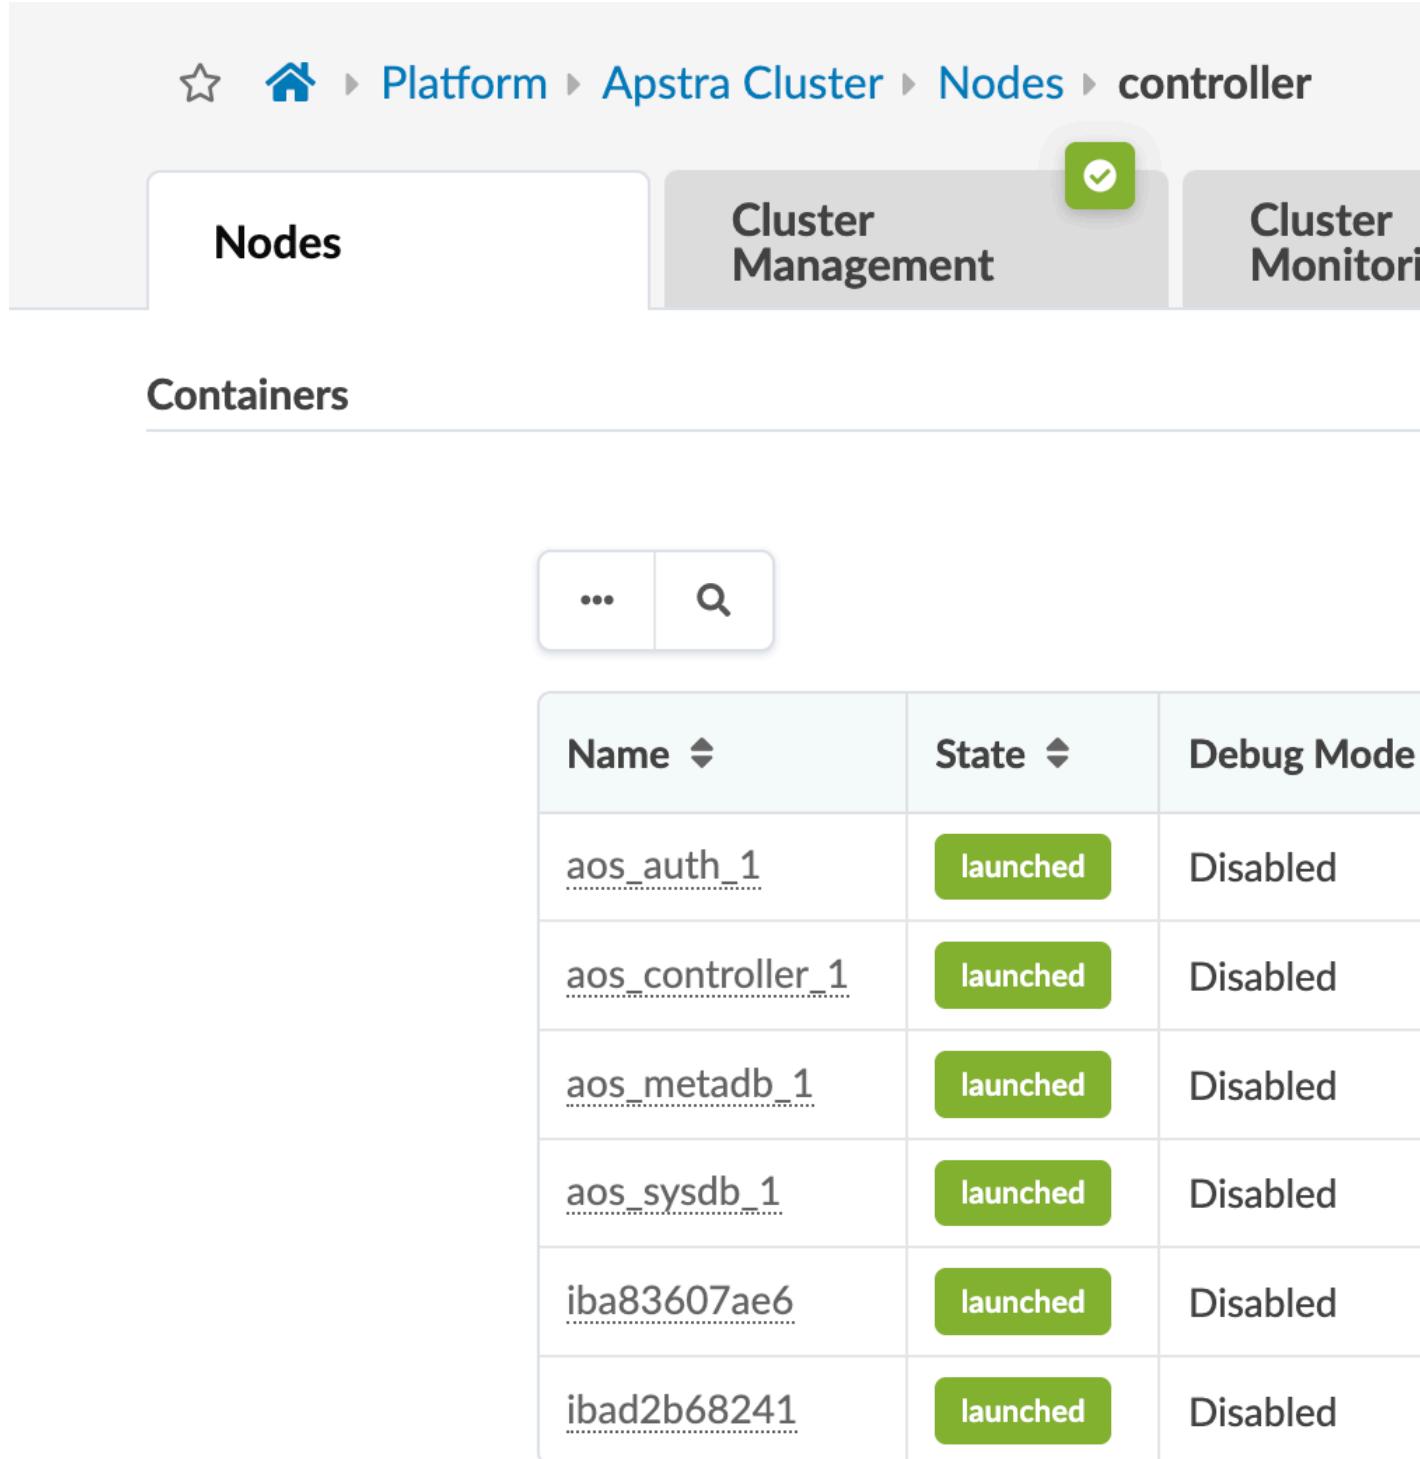
## Apstra Cluster Nodes Summary Table

The screenshot shows the Juniper Apstra Platform interface. The top navigation bar includes the Juniper Apstra logo, the version 5.0.0-64, and a star icon. The main navigation bar has a 'Platform' tab with a red badge showing '1', followed by an upward arrow icon. The left sidebar lists various platform management sections: User Management, Users, Roles, Security, Allowed List, Banned List, ACL, Ratelimit Configuration, Password Complexity ..., External Services, Syslog Configuration, Streaming, Receivers, Global Statistics, and Event Log. To the right, a 'Nodes' summary table is displayed with columns for Address and Name. The table contains three rows with the following data:

Address	Name
10.29.136.3	controller
10.29.136.4	worker-1
10.29.136.5	worker-2

## Apstra Cluster Controller Node

To see controller usage details, click the controller address in the summary table, and scroll down to the **Containers** section. Our example controller node includes 6 containers, as shown below.



The screenshot shows the Apstra Platform interface with the following details:

- Navigation Bar:** Platform > Apstra Cluster > Nodes > controller
- Tabs:** Nodes (selected), Cluster Management, Cluster Monitor
- Section:** Containers
- Search Bar:** ... (three dots) and a magnifying glass icon
- Table:** Displays 6 containers with the following data:

Name	State	Debug Mode
aos_auth_1	launched	Disabled
aos_controller_1	launched	Disabled
aos_metadb_1	launched	Disabled
aos_sysdb_1	launched	Disabled
iba83607ae6	launched	Disabled
ibad2b68241	launched	Disabled

## Apstra Cluster Worker Node

To see worker details, click one of the worker addresses in the summary table, and scroll down to the **Containers** section. Our example controller node includes 21 containers, one for each offbox agent.

☆  Platform > Apstra Cluster > Nodes > worker-1

Nodes Cluster Management Cluster Monitoring

Containers

...

Name State Debug

aos-offbox-10_21-f	launched	Not Supp
aos-offbox-10_23-f	launched	Not Supp
aos-offbox-10_24-f	launched	Not Supp
aos-offbox-10_31-f	launched	Not Supp
aos-offbox-10_32-f	launched	Not Supp
aos-offbox-10_37-f	launched	Not Supp
aos-offbox-10_39-f	launched	Not Supp
aos-offbox-10_40-f	launched	Not Supp

## Required Communication Ports

Open ports and services that run on the Apstra server are listed in the table below.

Apstra requires a minimum of eight (8) SSH connections, two (2) SSH max-sessions-per-connection, and twenty (20) SSH rate-limit (maximum number of connection attempts per minute).

A running iptables instance ensures that network traffic to and from the Apstra server is restricted to the services listed.

**Table 3: Apstra Server Network Protocol Requirements**

Source	Destination	Protocol	Description
User workstation	Apstra Server	tcp/22 (ssh)	CLI access to Apstra server
User workstation	Apstra Server	tcp/80 (http)	Redirects to tcp/443 (https)
User workstation	Apstra Server	tcp/443 (https)	GUI and REST API
Network Device for device agents	Apstra Server	tcp/80 (http)	Redirects to tcp/443 (https)
Network Device and Offbox Agent	Apstra Server	tcp/443 (https)	Device agent installation and upgrade, Rest API
Network Device or Offbox Agent	Apstra Server	tcp/29730-29739	Agent binary protocol (Sysdb)
ZTP Server	Apstra Server	tcp/443 (https)	Rest API for Device System Agent Install
Apstra Server	Network Devices	tcp/22 (ssh)	Device agent installation and upgrade
Offbox Agent	Network Devices	tcp/32767 (grpc/ssl)	Junos streaming telemetry using gRPC over SSL
Offbox Agent	Network Devices	tcp/443 (https) tcp/9443 (nxapi) tcp/830 (for Junos)	Management from Offbox Agent

## Additional Network Protocols

The network protocols in the table below are not required for Apstra server functionality, but they may be required for network device configuration and discovery, and for direct access to devices.

**Table 4: Additional Network Protocols**

Source	Destination	Protocol	Description
Administrator	Network Device	tcp/22 (ssh)	Device management from Administrator
Network Device	DNS Server	udp/53 (dns)	DNS Discovery for Apstra server IP (if applicable)
Network Device	DHCP Server	udp/67-68 (dhcp)	DHCP for automatic management IP (if applicable)
		(icmp type 0, type 8 for echo and response)	As necessary for network troubleshooting. Not required for the Apstra server.

## Network Client Services

Use and configuration of the Apstra server determine the number of network client services that must be enabled.

**Table 5: Apstra server Network Client Services**

Source	Destination	Protocol	Description
Apstra Server	DNS Server	udp/53 (dns)	Server DNS Client
Apstra Server	LDAP Server	tcp/389 (ldap) tcp/636 (ldaps)	Apstra Server LDAP Client (if configured)
Apstra Server	TACACS+ Server	tcp/udp/49 (tacacs)	Apstra Server TACACS+ Client (if configured)
Apstra Server	RADIUS Server	tcp/udp/1812 (radius)	Apstra Server RADIUS Client (if configured)
Apstra Server	Syslog Server	udp/514 (syslog)	Apstra Server Syslog Client (if configured)

## Install Apstra on ESXi

These instructions are for *installing Apstra software* on an ESXi hypervisor. For information about using ESXi in general, refer to VMware's [ESXi documentation](#).

1. Confirm that you're running one of the "[Supported Hypervisors and Versions](#)" on page 2.
2. Make sure that you allocate enough resources for your network. For example, if you're using more than 25 offbox agents, you'll need to add another VM. See "[Required Server Resources](#)" on page 2 for more details.



**CAUTION:** If you allocate insufficient CPU and RAM allocations for your network you'll probably experience issues such as slow response times, system errors and a critical "segmentation fault" (core dump). If this happens, you'll need to delete the VM and redeploy it with additional resources.

3. Apstra software is delivered pre-installed on a single VM. The same Apstra VM image is used for installing both the Apstra controller and Apstra workers. As a registered support user, download the [Apstra VM Image for VMware ESXi \(OVA\)](#) from [Juniper Support Downloads](#).



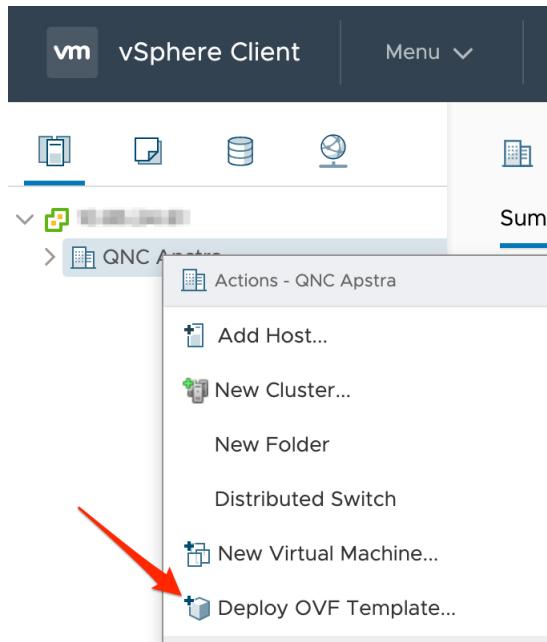
To download the image on your localhost, [CLICK HERE](#)

To download the image directly on your device, use the following URL:

[https://cdn.juniper.net/software/jafc/4.0.2-aos\\_server\\_4.0.2-142.ova?](https://cdn.juniper.net/software/jafc/4.0.2-aos_server_4.0.2-142.ova?)

[copy](#)

4. Log in to vCenter, right-click your target deployment environment, then click **Deploy OVF Template**.



5. Specify the URL or local file location for the OVA file you downloaded, then click **Next**.

## Deploy OVF Template

### 1 Select an OVF template

- 2 Select a name and folder
- 3 Select a compute resource
- 4 Review details
- 5 Select storage
- 6 Ready to complete

### Select an OVF template

Select an OVF template from remote URL or local file system

Enter a URL to download and install the OVF package from the location accessible from your computer, such as a local hard drive, CD/DVD drive.

URL

http | https://remoteserver-address/filetodeploy.ovf | .ova

Local file

**Choose Files** aos\_server\_4.0.2-142.ova

6. Specify a unique name and target location for the VM, then click **Next**.

## Deploy OVF Template

✓ 1 Select an OVF template

**2 Select a name and folder**

3 Select a compute resource

4 Review details

5 Select storage

6 Ready to complete

**Select a name and folder**

Specify a unique name and target location

Virtual machine name: aos\_server4.0.2-142

Select a location for the virtual machine.



7. Select your destination compute resource, then click **Next**.

## Deploy OVF Template

✓ 1 Select an OVF template

✓ 2 Select a name and folder

**3 Select a compute resource**

4 Review details

5 Select storage

**Select a compute resource**

Select the destination compute resource for this operation



8. Review template details, then click **Next**.
9. Select storage for the files, then click **Next**. We recommend thick provisioning for the Apstra server.

## Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- 5 Select storage**
- 6 Select networks
- 7 Ready to complete

### Select storage

Select the storage for the configuration and disk files

Encrypt this virtual machine (Requires Key Management Service)

Select virtual disk format:

Thick Provision L

VM Storage Policy:

Data

Name	Capacity	Provisioned
datastore1	215 GB	261.57 GB
datastore1 (6)	215 GB	493.67 GB
NFS-Datastore	2 TB	1.73 TB

10. Map the Apstra Management network to enable it to reach the virtual networks that the Apstra server will manage on ESXi, then click **Next**.

## Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 Select storage
- 6 Select networks**
- 7 Ready to complete

### Select networks

Select a destination network for each source network.

Source Network	Destination Network
VM Network	topology1

## IP Allocation Settings

IP allocation:

Static - Manual

IP protocol:

IPv4

11. Review your specifications, then click **Finish**.

You're ready to ["configure"](#) on page 25 the Apstra server.

## Install Apstra on KVM

### SUMMARY

You can install KVM with [Virtual Machine Manager](#) or with the CLI.

### IN THIS SECTION

-  [Install on KVM with Virtual Machine Manager | 14](#)
-  [Install on KVM with CLI | 18](#)

These instructions are for *installing Apstra software* on a KVM hypervisor. For information about using KVM in general, refer to Linux [KVM documentation](#).

### Install on KVM with Virtual Machine Manager

1. Confirm that you're running one of the "[Supported Hypervisors and Versions](#)" on page 2.
2. Make sure that you allocate enough resources for your network. For example, if you're using more than 25 offbox agents, you'll need to add another VM. See "[Required Server Resources](#)" on page 2 for more details.



**CAUTION:** If you allocate insufficient CPU and RAM allocations for your network you'll probably experience issues such as slow response times, system errors and a critical "segmentation fault" (core dump). If this happens, you'll need to delete the VM and redeploy it with additional resources.

3. Apstra software is delivered pre-installed on a single VM. The same Apstra VM image is used for installing both the Apstra controller and Apstra workers. As a registered support user, download the [Apstra VM Image for Linux KVM \(QCOW2\)](#) from [Juniper Support Downloads](#).



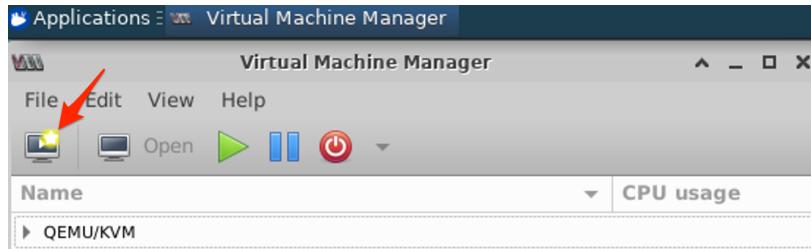
To download the image on your localhost, [CLICK HERE](#)

To download the image directly on your device, use the following URL:

[https://cdn.juniper.net/software/jafc/4.0.2-aos\\_server\\_4.0.2-142.qcow2.gz?](https://cdn.juniper.net/software/jafc/4.0.2-aos_server_4.0.2-142.qcow2.gz?)

[copy](#)

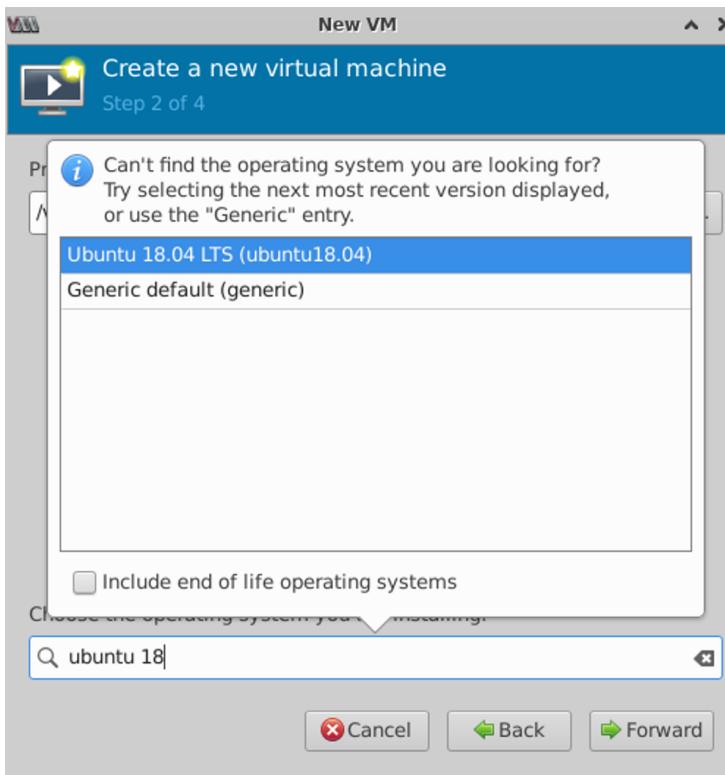
4. Uncompress the disk image, then move it to where it will run.
5. Start **Virtual Machine Manager**, then click the **Create a new virtual machine** button.



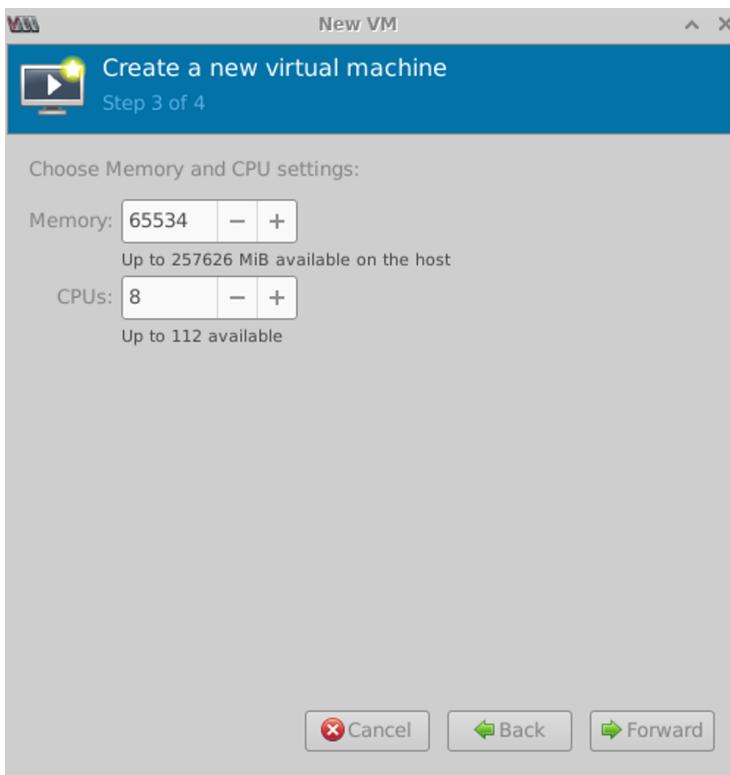
6. Select **Import existing disk image**, then click **Forward**.



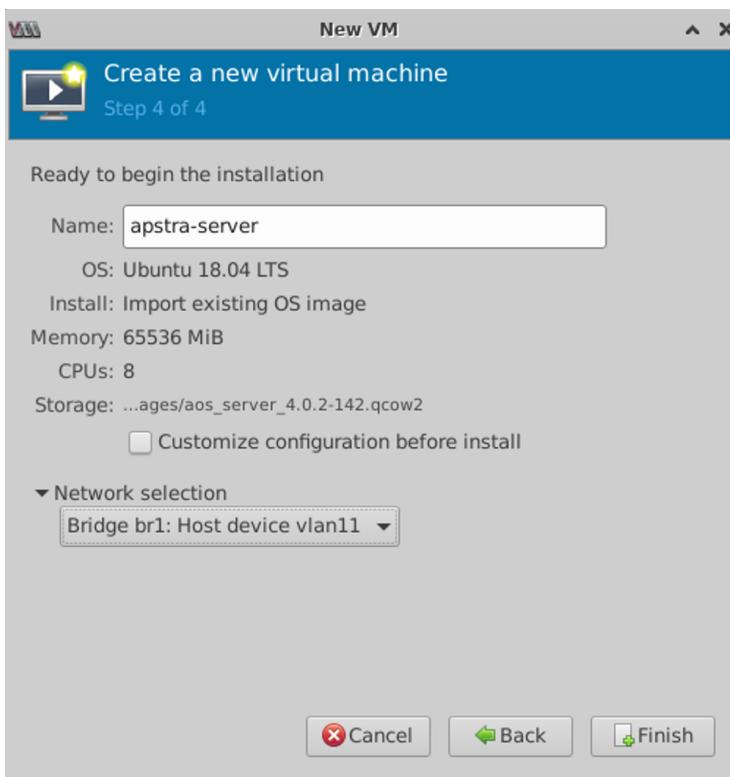
7. Browse to where you moved the QCOW2 image, then click **Choose Volume**.
8. Select **Ubuntu 18.04 LTS** operating system, then click **Forward**.



9. Specify memory and CPU requirements based on your environment.



10. Change the default name (optional), select the VM network that you want the VM to connect to, then click **Finish**. It may take a few minutes to create the VM.



You're ready to "configure" on page 25 the Apstra server.

## Install on KVM with CLI

1. Confirm that you're running one of the "Supported Hypervisors and Versions" on page 2.
2. Make sure that you allocate enough resources for your network. For example, if you're using more than 25 offbox agents, you'll need to add another VM. See "Required Server Resources" on page 2 for more details.



**CAUTION:** If you allocate insufficient CPU and RAM allocations for your network you'll probably experience issues such as slow response times, system errors and a critical "segmentation fault" (core dump). If this happens, you'll need to delete the VM and redeploy it with additional resources.

3. Ensure that the QEMU environment and bridge networking are installed and configured. For examples of installing and configuring QEMU, refer to the following documents:
  - Ubuntu - <https://help.ubuntu.com/community/KVM/Installation>
  - RHEL - [https://access.redhat.com/documentation/en-us/red\\_hat\\_enterprise\\_linux/7/html/virtualization\\_deployment\\_and\\_administration\\_guide/index](https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/virtualization_deployment_and_administration_guide/index)
4. You must use e1000 or virtio Linux KVM network drivers. Run the command `ethtool -i eth0` from the Apstra server to confirm which network drivers you're using.

```
admin@aos-server:~$ ethtool -i eth0
driver: virtio_net
version: 1.0.0
firmware-version:
expansion-rom-version:
bus-info: 0000:00:03.0
supports-statistics: no
supports-test: no
supports-eeprom-access: no
supports-register-dump: no
supports-priv-flags: no
admin@aos-server:~$
```



**CAUTION:** Using other drivers such as rtl8139 may result in high CPU utilization for the ksoftirqd process.

5. As a registered support user, download the **Apstra VM Image for Linux KVM (QCOW2)** from [Juniper Support Downloads](#).
6. Uncompress the disk image (with **gunzip**) and move it to where it will run.

```
ubuntu@ubuntu:~$ ls -l
total 1873748
-rw-r--r-- 1 ubuntu ubuntu 1918712115 Feb  4 22:28 aos_server_4.0.2-142.qcow2.gz
ubuntu@ubuntu:~$ gunzip aos_server_4.0.2-142.qcow2.gz
ubuntu@ubuntu:~$ ls -l
total 1905684
-rw-r--r-- 1 ubuntu ubuntu 1951413760 Feb  4 22:28 aos_server_4.0.2-142.qcow2.gz
ubuntu@ubuntu:~$
```

7. Create a VM with the **virt-install** command line tool. For example, to install the **aos\_server\_4.0.2-142.qcow2.gz** image using the existing bridge network (named **br0**), use the following command:

```
ubuntu@ubuntu:~$ sudo virt-install --name=aos-server --disk=aos_server_4.0.2-142.qcow2 --os-type=linux --os-variant ubuntu18.04 --import --noautoconsole --vcpu=8 --ram=65536 --network bridge=br0,model=virtio

Starting install...
Domain creation completed.
ubuntu@ubuntu:~$ sudo virsh list
  Id   Name           State
  --
  4    aos-server     running
ubuntu@ubuntu:~$
```

8. Connect to the VM console.

```
ubuntu@ubuntu:~$ sudo virsh console aos-server
Connected to domain aos-server
Escape character is ^]

Apstra Operating System (AOS)

aos-server login:
```

You're ready to "configure" on page 25 the Apstra server.

## Install Apstra on Hyper-V

These instructions are for *installing Apstra software* on a Microsoft Hyper-V hypervisor. For information about using Hyper-V in general, refer to Microsoft's [Hyper-V documentation](#).

1. Confirm that you're running one of the "[Supported Hypervisors and Versions](#)" on page 2.
2. Make sure that you allocate enough resources for your network. For example, if you're using more than 25 offbox agents, you'll need to add another VM. See "[Required Server Resources](#)" on page 2 for more details.



**CAUTION:** If you allocate insufficient CPU and RAM allocations for your network you'll probably experience issues such as slow response times, system errors and a critical "segmentation fault" (core dump). If this happens, you'll need to delete the VM and redeploy it with additional resources.

3. Apstra software is delivered pre-installed on a single VM. The same Apstra VM image is used for installing both the Apstra controller and Apstra workers. As a registered support user, download the [Apstra VM VHDX Image for Microsoft Hyper-V](#) from [Juniper Support Downloads](#).



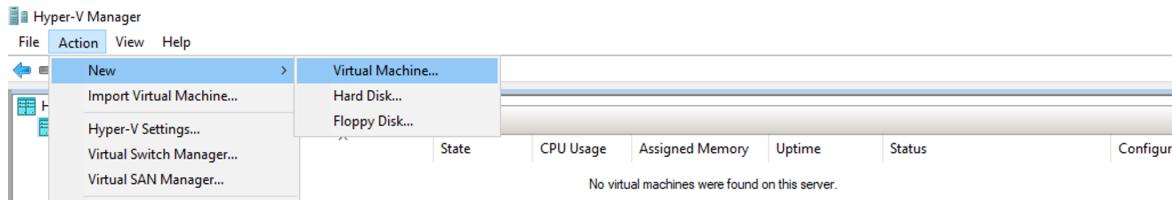
To download the image on your localhost, [CLICK HERE](#)

To download the image directly on your device, use the following URL:

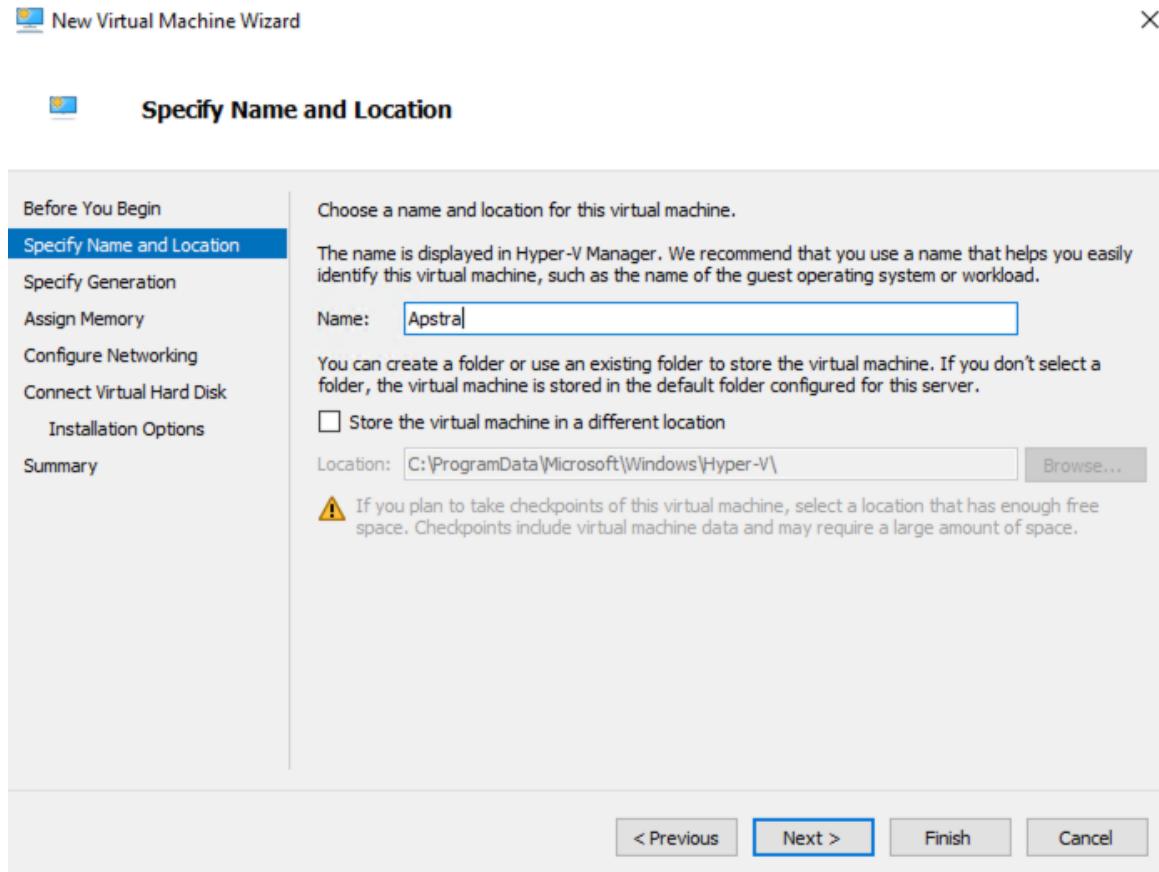
[https://cdn.juniper.net/software/jafc/4.0.1/aos\\_server\\_4.0.1-1045.vhdx.gz](https://cdn.juniper.net/software/jafc/4.0.1/aos_server_4.0.1-1045.vhdx.gz)

[copy](#)

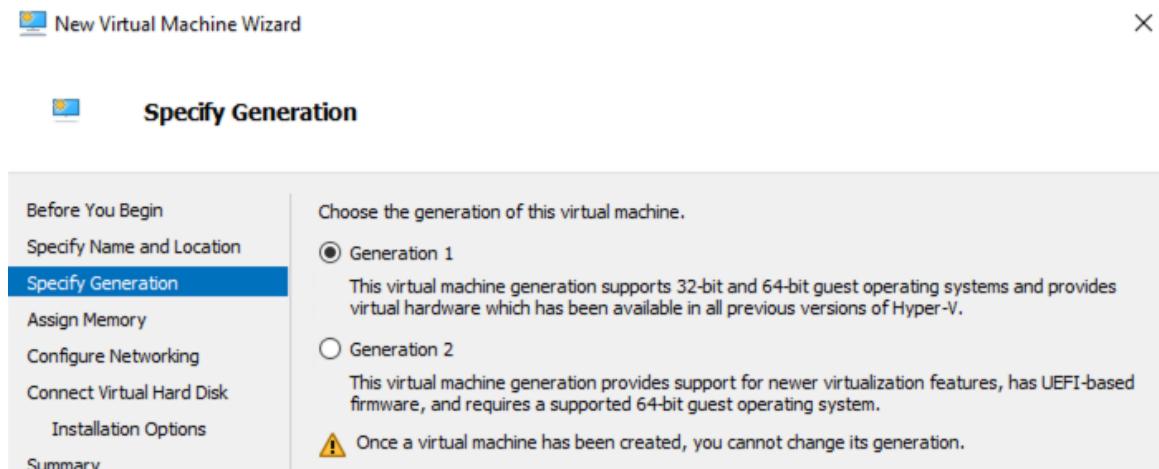
4. Uncompress the disk image and move it to where it will run.
5. Start Hyper-V Manager, select the server for the VM and navigate to **Actions > New > Virtual Machine**. The **New Virtual Machine Wizard** opens.



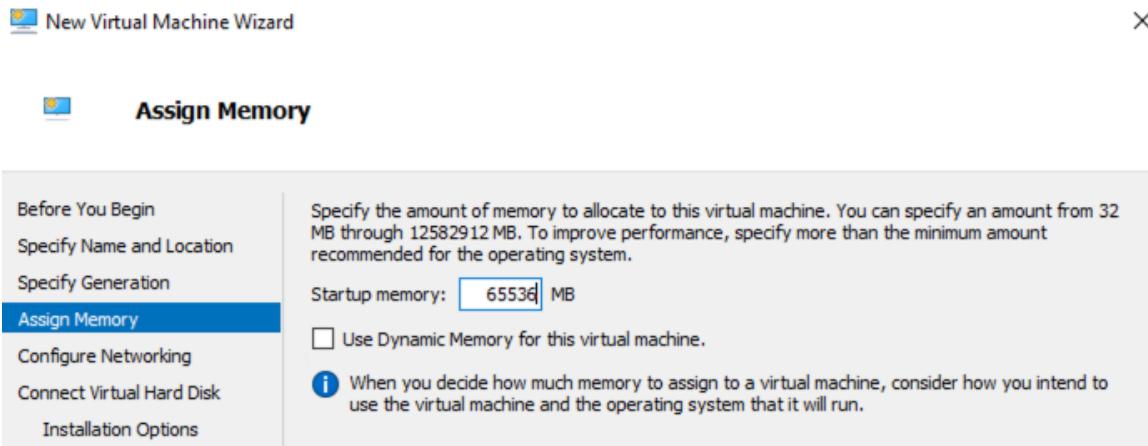
6. Specify a VM name and location, then click **Next**.



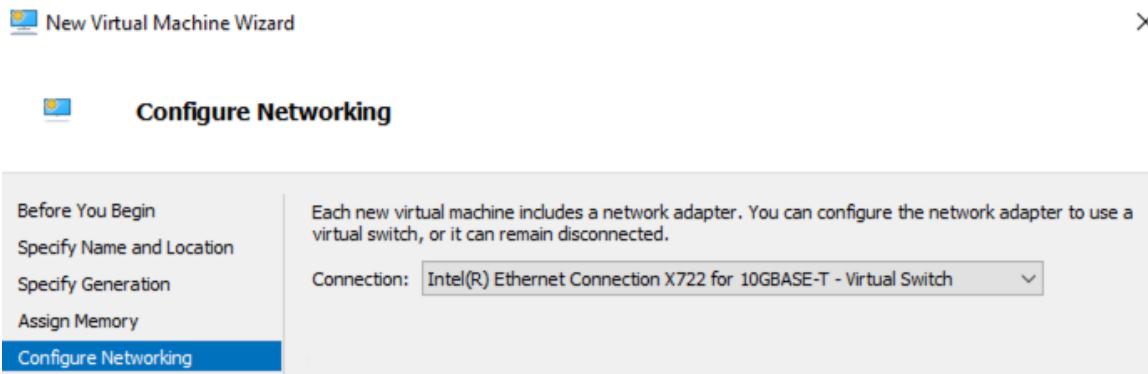
7. Specify **Generation 1**, then click **Next**.



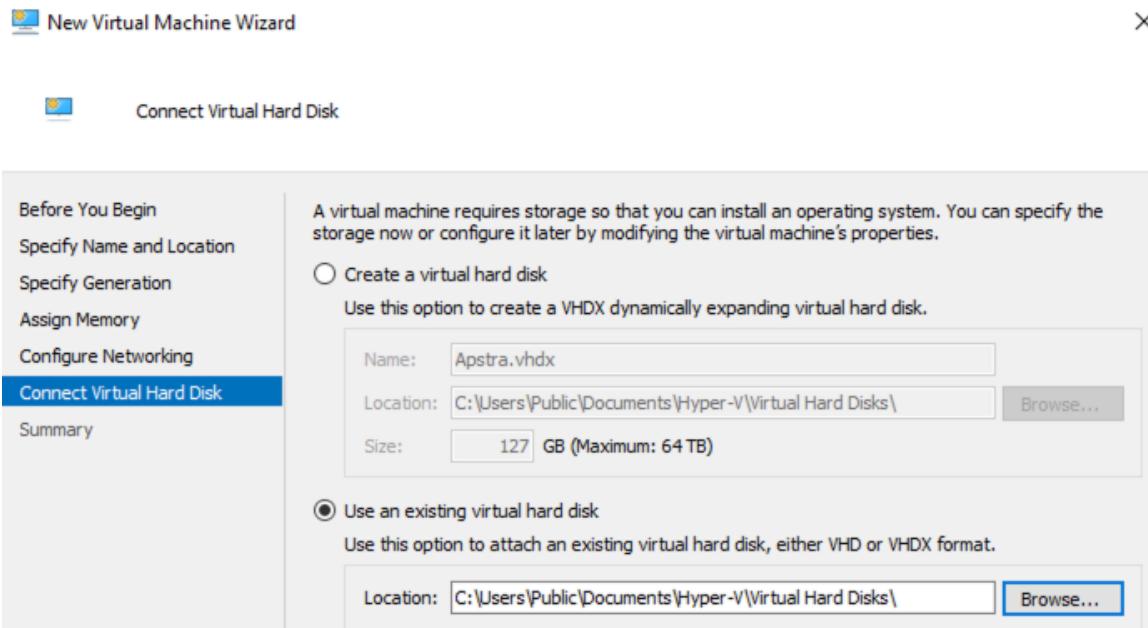
8. Specify required memory based on your environment, then click **Next**.



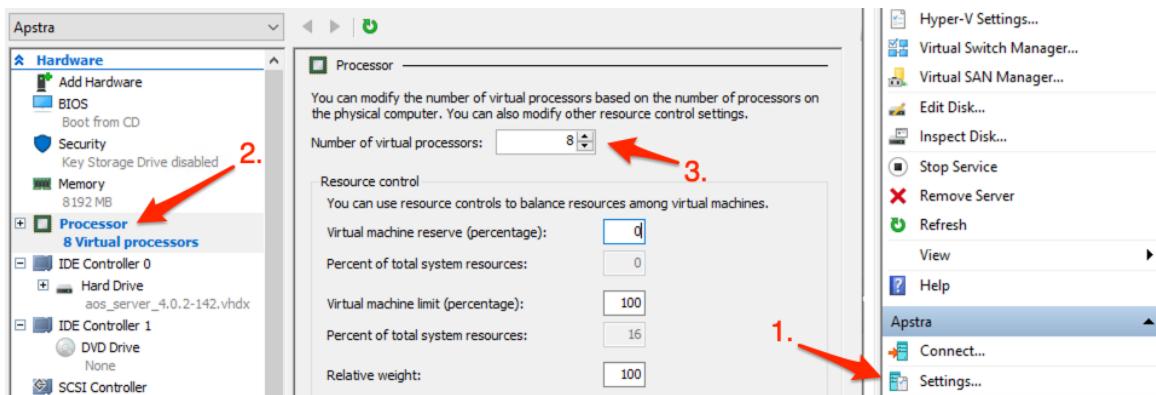
9. Configure the virtual switch as required for your deployment environment, then click **Next**.



10. Select **Use an existing virtual hard disk** and browse to the extracted file, then click **Finish**.



11. Click **Settings** (right panel), click **Processor** (left panel), specify the number of virtual processors based on required VM resources, then click **OK**.

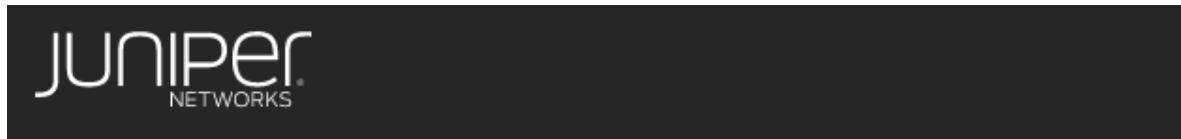


You're ready to "configure" on page 25 the Apstra server. (When the Apstra server is configured, the Docker daemon runs properly.)

## Install Apstra on VirtualBox

VirtualBox is for demonstration and lab purposes only. Production environments require a proper enterprise-scale virtualization solution (See "Supported Hypervisors and Versions" on page 2). These instructions are for *installing Apstra software* on a VirtualBox hypervisor. For information about using VirtualBox in general, refer to Oracle's [VirtualBox documentation](#) or the open-source community.

1. Apstra software is delivered pre-installed on a single virtual machine (VM). As a registered support user, download the **Apstra VM Image for VMware ESXi** (OVA) from [Juniper Support Downloads](#) to your local workstation.



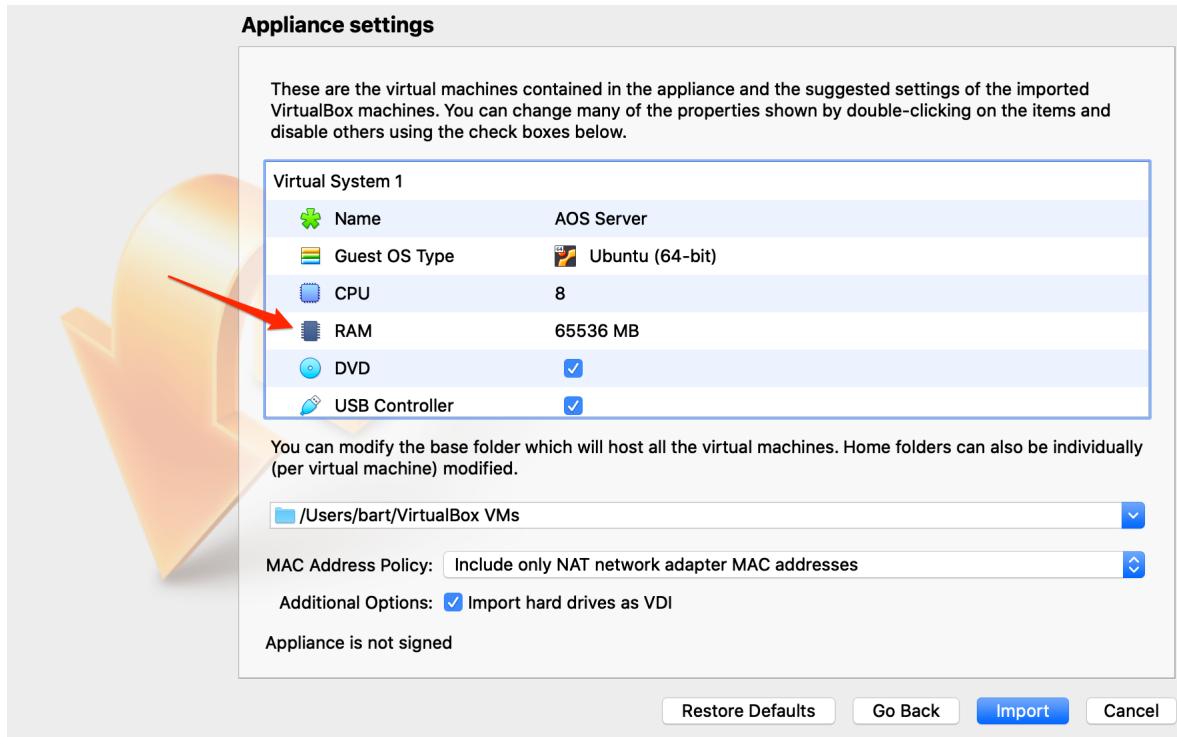
To download the image on your localhost, [CLICK HERE](#)

To download the image directly on your device, use the following URL:

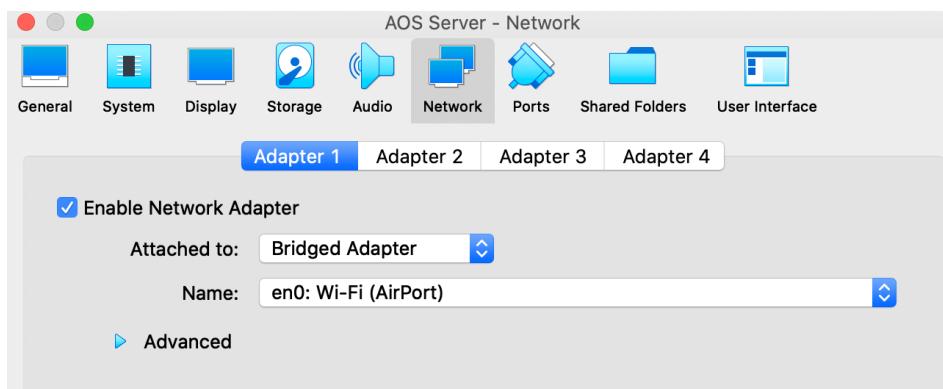
[https://cdn.juniper.net/software/jafc/4.0.2-aos\\_server\\_4.0.2-142.ova?](https://cdn.juniper.net/software/jafc/4.0.2-aos_server_4.0.2-142.ova?)

[copy](#)

2. Start VirtualBox, select **File > Import Appliance**, navigate to the OVA file, select it, then click **Continue**.
3. Change **RAM** to 8 GB. 8 GB is sufficient for lab and testing purposes.



4. Click **Import** to start the import process.
5. When the import is complete, start the server VM, click **Settings** and confirm that the VM meets requirements. In particular, check network settings for the adapter that's attached to your management network. If this value is not set correctly, the Apstra server does not receive an IP address. Since VirtualBox has one network adapter attached to the bridged adapter using active networking, full connectivity from your workstation to the VM (HTTP, SSH) is expected by default.



6. ["Configure" on page 25](#) the Apstra server.
7. Confirm connectivity. (Run ifconfig -a on the VM to get the IP address.)

- SSH from your workstation to the VM's active network adapter IP address.
- Point a web browser to the VM's active network adapter IP address.

## Configure Apstra Server

### IN THIS SECTION

- [Configure Apstra Server Static Management IP Address | 28](#)
- [Change Apstra Server Hostname | 28](#)
- [Configure Apstra Docker Subnets | 29](#)
- [Replace Apstra Server SSL Certificate | 30](#)

1. Log in to the Apstra server with the default credentials (user: **admin**, password: **admin**) either from the web console or via SSH (ssh admin@<apstra-server-ip> where <apstra-server-ip> is the IP address of the Apstra server.) You must change the default password before you can proceed.

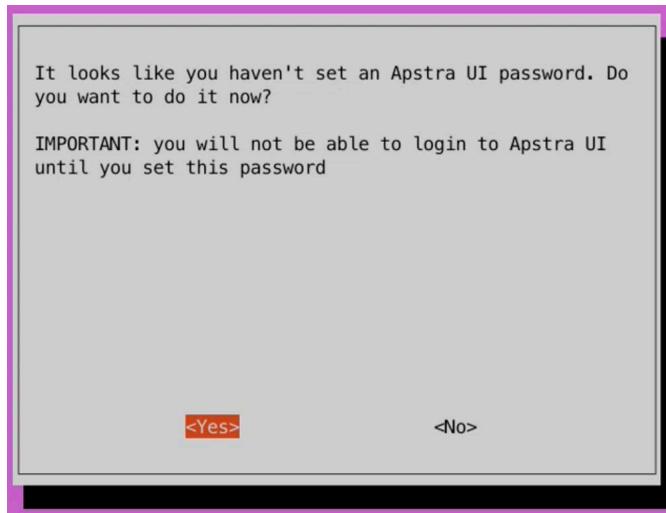
```
(admin@[REDACTED]) Password:  
(admin@[REDACTED]) You are required to change your password immediately (root  
Changing password for admin.  
(current) UNIX password:  
(admin@[REDACTED]) New password:
```

2. Enter a password that meets the following complexity requirements, then enter it again:

- Must contain at least 14 characters
- Must contain an uppercase letter
- Must contain a lowercase letter
- Must contain a digit
- Must contain a special character
- Must NOT be the same as the username
- Must NOT contain a repeat of the same character
- Must NOT contain consecutive sequential characters

- Must NOT use adjacent keys on the keyboard

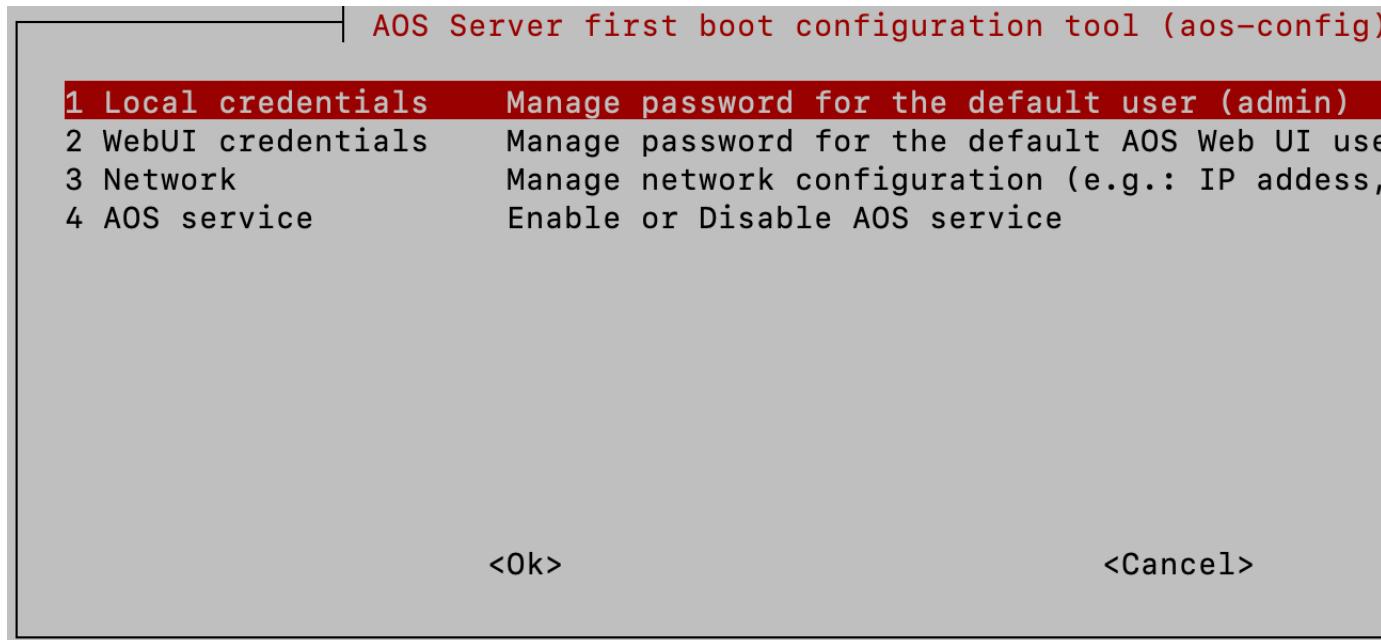
3. When you've successfully changed the Apstra server password a dialog opens prompting you to set the Apstra GUI password.



You won't be able to access the Apstra GUI until you set this password. Select **Yes** and enter a password that meets the following complexity requirements, then enter it again:

- Must contain at least 9 characters
- Must contain an uppercase letter
- Must contain a lowercase letter
- Must contain a digit
- Must contain a special character
- Must NOT be the same as the username
- Must NOT contain a repeat of the same character
- Must NOT contain consecutive sequential characters
- Must NOT use adjacent keys on the keyboard

4. A dialog appears stating "Success! Apstra UI password is changed." Select **OK**.  
5. The configuration tool menu appears.



You've just changed the local credentials and the Apstra GUI credentials, so you don't need to manage them again now.

The network is configured to use DHCP by default. To assign static IP addresses instead, select **Network**, change it to **Manual**, and provide the following:

- (Static Management) IP address in CIDR format with netmask (for example, 192.168.0.10/24)
- Gateway IP address
- Primary DNS
- Secondary DNS (optional)
- Domain

6. Apstra service is stopped by default. To start and stop Apstra service, select **AOS service** and select **Start** or **Stop**, as appropriate. Starting service from this configuration tool invokes `/etc/init.d-aos`, which is the equivalent of running the command `service aos start`.

7. To exit the configuration tool and return to the CLI, select **Cancel** from the main menu. (To open this tool again in the future, run the command `aos_config`.)

You're ready to "Replace the default SSL certificate with a signed one" on page 31.



**CAUTION:** We recommend that you back up the Apstra server on a regular basis (since HA is not available). For backup details, see the Apstra Server Management section of the Juniper Apstra User Guide. For information

about setting up automated backup collection see the [Juniper Support Knowledge Base article KB37808](#).

## Configure Apstra Server Static Management IP Address

If you're not using DHCP, you can use the ["configuration tool" on page 25](#) to enter a static management IP address:

1. Log in to the Apstra server via console, as user **admin**.
2. Run the configuration tool `aos_config` and configure Network to set a static management IP address.
3. If you have already installed onbox agents, you must reconfigure each device agent (`/mnt/flash/aos-config`, `/etc/aos/aos.conf`) to point to the new Apstra server IP address.

## Change Apstra Server Hostname

### IN THIS SECTION

- [Configure New SSH Keys on Apstra Server | 28](#)



**CAUTION:** To avoid issues with the Apstra container's binding, don't change the `/etc/hostname` file directly with any Linux CLI command or other command than the one below.

1. SSH into the Apstra server as user **admin**. (`ssh admin@<apstra-server-ip>` where `<apstra-server-ip>` is the IP address of the Astra server.)
2. With root privileges, run the command `/#aos_hostname <hostname>` where `<hostname>` is the new hostname of the Apstra server. This command modifies the hostname in the `/etc/hostname` file and performs necessary backend configuration.
3. For the change to take effect, reboot the Apstra server, preferably during a maintenance window. The Apstra server is temporarily unavailable during a reboot, though it most likely won't impact service.

## Configure New SSH Keys on Apstra Server

You can replace SSH host keys on new or existing Apstra server VMs.

1. SSH into the Apstra server as user **admin**. (`ssh admin@<apstra-server-ip>` where `<apstra-server-ip>` is the IP address of the Astra server.)
2. Run the command `sudo rm /etc/ssh/ssh_host*` to remove SSH host keys.

- Run the command `sudo dpkg-reconfigure openssh-server` to configure new SSH host keys.

```
admin@aos-server:~$ sudo dpkg-reconfigure openssh-server
Creating SSH2 RSA key; this may take some time ...
2048 SHA256:EWRFcs4V6Bm0ILR3T2Pxng1uE0qXQ/z9IKkXrnLpJs root@aos-server (RSA)
Creating SSH2 ECDSA key; this may take some time ...
256 SHA256:THaXEia8VW6Jfw60BXFegu1Cav0zcGSV0y9RkN0Pxf4 root@aos-server (ECDSA)
Creating SSH2 ED25519 key; this may take some time ...
256 SHA256:0H0n0nnF+7oRaF5HggI4vWeyxT+JNsHcbvNpBJdaKhQ root@aos-server (ED25519)
admin@aos-server:~$ sudo dpkg-reconfigure openssh-server
```

- To restart the SSH server process, run the command `sudo systemctl restart ssh`.

## Configure Apstra Docker Subnets

### IN THIS SECTION

- [docker0 | 29](#)
- [Docker Network and Apstra Upgrades | 30](#)

The Apstra server Docker containers require one network for internal connectivity, which is automatically configured with the following subnets:

- `docker0: inet 172.17.0.1/16`
- `Apstra same VM (in-place) upgrade docker network: inet 172.18.0.1/16`

If you need to use these subnets elsewhere, to avoid conflicts, change the Docker network as follows:

### **docker0**

Update `bip` with the new subnet. If the `/etc/docker/daemon.json` file doesn't already exist, create one with the following format (Replace `172.26.0.1/16` in the example below with your own subnet.):

```
$ sudo vi /etc/docker/daemon.json

{
  "bip": "172.26.0.1/16"
}
```

```
$ sudo service docker restart
$ sudo service aos restart
```

## Docker Network and Apstra Upgrades

To use a different subnet, create or edit the `/etc/docker/daemon.json` file with the following format (Replace `172.27.0.0/16` in the example with your own subnet).

```
$ sudo vi /etc/docker/daemon.json

{
  "default-address-pools":
  [
    {
      "base": "172.27.0.0/16",
      "size": 24
    }
  ]
}

$ sudo service docker restart
$ sudo service aos restart
```

## Replace Apstra Server SSL Certificate

### IN THIS SECTION

- [Replace SSL Certificate on Apstra Server with Signed One | 31](#)
- [Replace SSL Certificate on Apstra Server with Self-Signed One | 33](#)

For security, we recommend that you replace the default self-signed SSL certificate with one from your own certificate authority. Web server certificate management is the responsibility of the end user. Juniper support is best effort only.

## Replace SSL Certificate on Apstra Server with Signed One

When you boot up the Apstra server for the first time, a unique self-signed certificate is automatically generated and stored on the Apstra server at /etc-aos/nginx.conf.d (nginx.crt is the public key for the webserver and nginx.key is the private key.) The certificate is used for encrypting the Apstra server and REST API. It's not for any internal device-server connectivity. Since the HTTPS certificate is not retained when you back up the system, you must manually back up the etc-aos folder. We recommend replacing the default SSL certificate. Web server certificate management is the responsibility of the end user. Juniper support is best effort only.

1. Back up the existing OpenSSL keys.

```
admin@aos-server:/$ sudo -s
[sudo] password for admin:

root@aos-server:/# cd /etc-aos/nginx.conf.d
root@aos-server:/etc-aos/nginx.conf.d# cp nginx.crt nginx.crt.old
root@aos-server:/etc-aos/nginx.conf.d# cp nginx.key nginx.key.old
```

2. Create a new OpenSSL private key with the built-in openssl command.

```
root@aos-server:/etc-aos/nginx.conf.d# openssl genrsa -out nginx.key 2048
Generating RSA private key, 2048 bit long modulus
.....+++
.....+++
e is 65537 (0x10001)
```



**CAUTION:** Don't modify nginx.crt or nginx.key filenames. They're referred to in nginx.conf.

As part of subsequent service upgrades, these files could be replaced, so the filenames must be predictable.

Also, don't change configuration in nginx.conf, as this file may be replaced during Apstra server upgrade, and any changes you make would be discarded.

3. Create a certificate signing request. If you want to create a signed SSL certificate with a Subjective Alternative Name (SAN) for your Apstra server HTTPS service, you must manually create an OpenSSL template. For details, see [Juniper Support Knowledge Base article KB37299](#).



**CAUTION:** If you have created custom OpenSSL configuration files for advanced certificate requests, don't leave them in the Nginx configuration folder. On startup, Nginx will attempt to load them (\*.conf), causing a service failure.

```
root@aos-server:/etc/aos/nginx.conf.d# openssl req -new -sha256 -key nginx.key -out nginx.csr
You are about to be asked to enter information that will be incorporated
into your certificate request.

What you are about to enter is what is called a Distinguished Name or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.

-----
Country Name (2 letter code) [AU]:US
State or Province Name (full name) [Some-State]:California
Locality Name (eg, city) []:Menlo Park
Organization Name (eg, company) [Internet Widgits Pty Ltd]:Apstra, Inc
Organizational Unit Name (eg, section) []:
Common Name (e.g. server FQDN or YOUR name) []:aos-server.apstra.com
Email Address []:support@apstra.com

Please enter the following 'extra' attributes
to be sent with your certificate request
A challenge password []:
An optional company name []:
```

4. Submit your Certificate Signing Request (nginx.csr) to your Certificate Authority. The required steps are outside the scope of this document; CA instructions differ per implementation. Any valid SSL certificate will work. The example below is for self-signing the certificate.

```
root@aos-server:/etc/aos/nginx.conf.d# openssl req -x509 -sha256 -days 3650 -key nginx.key -
in nginx.csr -out nginx.crt
root@aos-server:/etc/aos/nginx.conf.d#
```

5. Verify that the SSL certificates match: private key, public key, and CSR.

```
root@aos-server:/etc/aos/nginx.conf.d# openssl rsa -noout -modulus -in nginx.key | openssl md5
(stdin)= 60ac4532a708c98d70fee0dbcaab1e75
```

```
root@aos-server:/etc/aos/nginx.conf.d# openssl req -noout -modulus -in nginx.csr | openssl md5
```

```
(stdin)= 60ac4532a708c98d70fee0dbcaab1e75

root@aos-server:/etc/aos/nginx.conf.d# openssl x509 -noout -modulus -in nginx.crt | openssl
md5
(stdin)= 60ac4532a708c98d70fee0dbcaab1e75
```

6. To load the new certificate, restart the nginx container.

```
root@aos-server:/etc/aos/nginx.conf.d# docker restart aos_nginx_1
aos_nginx_1
root@aos-server:/etc/aos/nginx.conf.d
```

7. Confirm that the new certificate is in your web browser and that the new certificate common name matches 'aos-server.apstra.com'.

### Replace SSL Certificate on Apstra Server with Self-Signed One

When you boot up the Apstra server for the first time, a unique self-signed certificate is automatically generated and stored on the Apstra server at /etc/aos/nginx.conf.d (nginx.crt is the public key for the webserver and nginx.key is the private key.) The certificate is used for encrypting the Apstra server and REST API. It's not for any internal device-server connectivity. Since the HTTPS certificate is not retained when you back up the system, you must manually back up the etc/aos folder. We support and recommend replacing the default SSL certificate.

1. Back up the existing OpenSSL keys.

```
admin@aos-server:/$ sudo -s
[sudo] password for admin:

root@aos-server:/# cd /etc/aos/nginx.conf.d
root@aos-server:/etc/aos/nginx.conf.d# cp nginx.crt nginx.crt.old
root@aos-server:/etc/aos/nginx.conf.d# cp nginx.key nginx.key.old
```

2. If a Random Number Generator seed file .rnd doesn't exist in /home/admin, create one.

```
root@aos-server:~# touch /home/admin/.rnd
root@aos-server:~#
```

**3. Generate a new OpenSSL private key and self-signed certificate.**

```
root@aos-server:/etc/aos/nginx.conf.d# openssl req -newkey rsa:2048 -nodes -keyout nginx.key -x509 -days 824 -out nginx.crt -addext extendedKeyUsage=serverAuth -addext subjectAltName=DNS:apstra.com
Generating a RSA private key
.....+++++
.....+++++
writing new private key to 'nginx.key'
-----
You are about to be asked to enter information that will be incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished Name or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) [AU]:US
State or Province Name (full name) [Some-State]:California
Locality Name (eg, city) []:Menlo Park
Organization Name (eg, company) [Internet Widgits Pty Ltd]:Apstra, Inc
Organizational Unit Name (eg, section) []:
Common Name (e.g. server FQDN or YOUR name) []:aos-server.apstra.com
Email Address []:support@apstra.com
root@aos-server:/etc/aos/nginx.conf.d#
```

**4. To load the new certificate, restart the nginx container.**

```
root@aos-server:/etc/aos/nginx.conf.d# docker restart aos_nginx_1
aos_nginx_1
root@aos-server:/etc/aos/nginx.conf.d
```

# Apstra Upgrade

## IN THIS SECTION

- [Supported Apstra Server Upgrade Paths | 35](#)
- [Apstra Server Upgrade Workflow | 36](#)
- [Upgrade Apstra on New VM \(VM-VM\) | 38](#)
- [Roll Back Apstra Server Upgrade | 50](#)

## Supported Apstra Server Upgrade Paths

### IN THIS SECTION

- [Check Apstra Version | 35](#)
- [Upgrade to Version 5.0.0/5.0.1 | 36](#)

### Check Apstra Version

- The Apstra version is displayed in the GUI in the upper left corner of the left navigation menu under the Juniper Apstra logo. You can also check the version from the left navigation menu by navigating to **Platform > About**.
- To check your current Apstra version from the CLI, run the command `service aos show_version`.

## Upgrade to Version 5.0.0/5.0.1

**Table 6: Apstra Server Supported Upgrade Paths to Version 5.0.0/5.0.1**

From Version	VM-to-VM (on New VM)	In-Place (on Same VM)
4.2.2	Yes	No
4.2.1	Yes	No
4.2.0	Yes	No

## Apstra Server Upgrade Workflow

Stage	Description
Pre-Upgrade Validation	<ul style="list-style-type: none"> <li>Verify upgrade path is supported</li> <li>Verify sufficient VM memory for new version</li> <li>Check the new Apstra version release notes for config-rendering changes that could impact the data plane. Update configlets, as needed.</li> <li>Review blueprints and address issues</li> <li>Verify device models and OS versions are supported.</li> <li>If you're using Junos devices, make sure pristine configuration includes <code>mgmt_junos</code> VRF.</li> <li>Remove any device AAA configuration</li> <li>Remove any configlets used to configure firewalls</li> <li>Back up current Apstra environment</li> </ul>

*(Continued)*

Stage	Description
Upgrade Apstra Server	<ul style="list-style-type: none"> <li>Download Apstra VM image</li> <li>Install software on controller VM (Check configlets for conflicts.)</li> <li>Install software on worker VMs (new VMs with Apstra Cluster only)</li> <li>Verify connections to new server</li> <li>Import SysDB database</li> <li>Verify that configlets don't conflict with newly rendered config</li> <li>Log in to the upgraded server</li> <li>Change operation mode to normal</li> </ul>
Upgrade Agents	<ul style="list-style-type: none"> <li>From the Apstra GUI</li> </ul>
Upgrade Worker Nodes (Apstra Cluster only)	<ul style="list-style-type: none"> <li>If using new VMs for worker nodes - import state from new controller</li> <li>If reusing VMs for worker nodes - install software on worker VMs(s)</li> </ul>
Upgrade Device NOS, as needed	<p>If the NOS versions of your devices are not qualified on the new Apstra version, upgrade them to a qualified version. (See the <a href="#">Juniper Apstra User Guide</a> for details.)</p>
Shut down old Apstra server	<p>If you're upgrading an Apstra cluster and you replaced your worker nodes with new VMs, shut down the old worker VMs as well.</p>
Roll back Apstra Server, as needed	<p>You can roll back to a previous Apstra version, if needed.</p>

## Upgrade Apstra on New VM (VM-VM)

### IN THIS SECTION

- [Step 1: Pre-Upgrade Validation | 38](#)
- [Step 2: Deploy New Apstra Server | 40](#)
- [Step 3: Import State | 41](#)
- [Step 4: Import State for Intent-Based Analytics | 45](#)
- [Step 5: Keep Old VM's IP Address \(Optional\) | 45](#)
- [Step 6: Change Operation Mode to Normal | 46](#)
- [Step 7: Upgrade Onbox Agents | 48](#)
- [Step 8: Shut Down Old Apstra Server | 49](#)

By upgrading Apstra on a new VM, you'll receive Ubuntu Linux OS fixes, including security vulnerability updates. To upgrade the Apstra server you need Apstra OS admin user privileges and Apstra admin user group permissions.

### Step 1: Pre-Upgrade Validation

1. Refer to "[Upgrade Paths](#)" on page 35 to confirm that you're upgrading to a supported version. To find your current Apstra version in the Apstra GUI, look in the upper left corner of the left navigation menu under the Juniper Apstra logo.
2. Log in to the Apstra server as **admin** (For example, if your Apstra server IP address were 10.28.105.3, the command would be `ssh admin@10.28.105.3`).
3. Run the command `service aos status` to check that the server is active and has no issues.

```
admin@aos-server:~$ service aos status
* aos.service - LSB: Start AOS management system
  Loaded: loaded (/etc/init.d/aos; generated)
  Active: active (exited) since Thu 2022-10-28 17:11:27 UTC; 24h ago
    Docs: man:systemd-sysv-generator(8)
  Process: 1157 ExecStart=/etc/init.d/aos start (code=exited, status=0/SUCCESS)
```

4. Check the new Apstra version [release notes](#) for configuration-rendering changes that could impact the data plane.

- Review each blueprint to confirm that all **Service Config** is in the SUCCEEDED state. If necessary, undeploy and remove devices from the blueprint to resolve any pending or failed service config.

Category	Status	Count
Service Config	SUCCEEDED	4
Service Config	PENDING	0
Service Config	FAILED	0
Discovery Config	SUCCEEDED	0
Discovery Config	PENDING	0
Discovery Config	FAILED	0
Drain Config	SUCCEEDED	0
Drain Config	PENDING	0
Drain Config	FAILED	0

- Review each blueprint for probe anomalies, and resolve them as much as possible. Take notes of any remaining anomalies.
- Refer to the [Apstra User Guide](#) (References > Devices > Qualified Devices and NOS Versions) to verify that the device models and NOS versions are qualified on the new Apstra version. Upgrade or downgrade as needed to one of the supported versions.
- If you're using Junos devices, the pristine configuration must include the essential `mgmt_junos` VRF. See [Juniper Support Knowledge Base article KB77094](#) for more information.



**CAUTION:** If the pristine configuration doesn't include the `mgmt_junos` VRF, then deployment will fail.

- Remove any Device AAA configuration. During device upgrade, configured device agent credentials are required for SSH access.
- Remove any configlets used to configure firewalls. If you use FW's Routing Engine filters on devices, you'll need to update them to include the IP address of the new controller and worker VMs.
- To upgrade device system agents, Apstra must be able to SSH to all devices using the credentials that were configured when creating the agents. To check this from the Apstra GUI, navigate to **Devices > Managed Devices**, select the check box(es) for the device(s) to check, then click the **Check** button in the Agent menu. Verify that the states of all jobs is **SUCCESS**. If any check job fails, resolve the issue before proceeding with the Apstra upgrade.
- As **root** user, run the command `sudo aos_backup` to back up the Apstra server.

```
admin@aos-server:~$ sudo aos_backup
[sudo] password for admin:
=====
Backup operation completed successfully.
=====
```

```
New AOS snapshot: 2022-10-29_18-58-56
admin@aos-server:~$
```



**CAUTION:** The upgraded Apstra server doesn't include any time voyager revisions, so if you need to revert back to a past state, this backup is required. Previous states are not included due to the tight coupling with the reference designs which may change between Apstra versions.

13. Copy the backup files from `/var/lib-aos/snapshot/<snapshot_name>` to an external location.
14. Make sure that the new VM has the "[Required Server Resources](#)" on page 2 for the Apstra server.

## Step 2: Deploy New Apstra Server



**NOTE:** If you customized the `/etc-aos-aos.conf` file in the old Apstra server (for example, if you updated the `metadb` field to use a different network interface), you must re-apply the changes to the same file in the new Apstra server VM. It's not migrated automatically.

1. As a registered support user, [download the Apstra VM image from Juniper Support Downloads](#) (for example, `aos_server_5.0.0-63`) and transfer it to the new Apstra server.
2. Install and configure the new Apstra VM image with the new IP address (same or new FQDN may be used).
3. If you're using an Apstra cluster (offbox agents, IBA probes) and you're going to reuse your worker VMs, install the new software by running `sudo bash aos_<aos_version>.run`. If you're using *new* worker VMs, skip this step.



**NOTE:** Example of replacing all VMs: if you have a controller and 2 worker nodes and you want to upgrade all of them to new VMs, you would create 3 VMs with the new Apstra version and designate one of them to be the controller.

4. Verify that the new Apstra server has SSH access to the old Apstra server.
5. Verify that the new Apstra server can reach system agents. (See "[Required Communication Ports](#)" on page 8.)
6. Verify that the new Apstra server can reach applicable external systems (such as NTP, DNS, vSphere server, LDAP/TACACs+ server and so on).

### Step 3: Import State



**CAUTION:** If you perform any API/GUI write operations to the old Apstra server after you've started importing the new VM, those changes won't be copied to the new Apstra server.

1. Log in to the new Apstra server as user **admin**.
2. Run the `sudo aos_import_state` command to import SysDB from the old server, apply necessary translations, and import configuration. Include the following arguments, as applicable:
  - `--ip-address <old-apstra-server-ip>`
3. Enter the admin username:
  - `--username <admin-username>`
  - For Apstra clusters with new worker node IP addresses, include the following: `--cluster-node-address-mapping <old-node-ip> <new-node-ip>`
  - To keep the old Apstra server, add the `keep` argument to the `--disable-original-apstra-server {prompt, disable, keep}` command.

For example: `--disable-original-apstra-server keep`

- To run the upgrade preconditions checks without running the actual upgrade use the following: `--dry-run-connectivity-validation`
- To not check connectivity validation include the following: `--skip-connectivity-validation`
- If SSH credentials on your older Apstra version are not as strict as the requirements for the new Apstra version, then you need to add the `--override-cluster-node-credentials` argument to the `aos_import_state` command when importing your database to the new Apstra version. Otherwise, the upgrade will fail.
- To perform a no fail upgrade process when Junos devices are present without any configured management VRFs, add the `--skip-junos-mgmt-vrf-check` argument.

Example command: Single VM or Apstra Cluster with Same Worker Nodes

```
admin@aos-server:~$ sudo aos_import_state --ip-address 10.28.105.3 --username admin
```

### Example Command: Apstra Cluster with New Worker Nodes

```
admin@aos-server:~$ sudo aos_import_state --ip-address 10.28.105.3 --username admin --cluster-node-address-mapping 10.28.105.4 10.28.105.6 --cluster-node-address-mapping 10.28.105.7 10.28.105.8
```

In the example above, 10.28.105.4 and 10.28.105.7 are old worker node IP addresses; 10.28.105.6 and 10.28.105.8 are new worker node IP addresses.

```
admin@aos-server:~$ sudo aos_import_state --ip-address 10.28.105.3 --username admin --cluster-node-address-mapping 10.28.105.4 10.28.105.6
[sudo] password for admin:
AOS[2022-10-27_20:17:23]: Initiating docker library import DONE
SSH password for remote AOS VM:
Root password for remote AOS VM:
AOS[2022-10-27_20:17:50]: Preparing to retrieve data from remote AOS Server. DONE
AOS[2022-10-27_20:18:29]: Retrieving data from remote AOS Server. This step can take up to 10 minutes DONE
AOS[2022-10-27_20:21:44]: Importing retrieved state to AOS. This step can take up to 30 minutes DONE
AOS[2022-10-27_20:21:48]: Waiting for blueprint <3db44826-807f-4ab9-8ca0-e25040af7ef6> processing to finish. DONE
AOS[2022-10-27_20:21:55]: Waiting for blueprint <964211f7-7f3c-4b0a-b6b7-137790c461f5> processing to finish. DONE
Summary saved to /tmp-aos-upgrade-config-summary-2022.10.27-202203
```

Root is required for importing the database, so you'll be asked for the SSH password and root password for the remote Apstra VM.



**NOTE:** When you upgrade an Apstra cluster, the SSH password for old controller, old worker and new worker must be identical, otherwise the upgrade fails authentication. In the above example, the password you enter for 'SSH password for remote AOS VM' is used for remote controller, old worker, and new worker VMs. (AOS-27351)

If you change the worker VMs' SSH password after the upgrade, then you also need to update the worker's password in the Apstra GUI (Platform > Apstra Cluster > Nodes).



**NOTE:** The size of the blueprint and the Apstra server VM resources determine how long it takes to complete the import. If the database import exceeds the default value (40 min or 2400 seconds), the operation may 'time out'. If this happens, you can increase the timeout value with the AOS\_UPGRADE\_DOCKER\_EXEC\_TIMEOUT command.

For example, the following command increases the time before timeout to 2 hours (7200 seconds).

```
admin@aos-server:~$ sudo AOS_UPGRADE_DOCKER_EXEC_TIMEOUT=7200 aos_import_state --ip-address 10.10.10.10 --username admin
```

The upgrade script presents a summary view of the devices within the fabric that will receive configuration changes during the upgrade. A warning appears on the screen recommending that you read [Release Notes](#) and ["Upgrade Paths" on page 35](#) documentation before proceeding. The release notes include a category for **Configuration Rendering Changes**. Configuration rendering changes are clearly documented at the top explaining the impact of each change on the network.

```
Apstra Upgrade Summary
=====
This is a summary of configuration pushed to devices logically grouped
into sections. Use 'q' to exit this view. For more device specific
configurations, use the menu after quitting this view

WARNING: This upgrade will modify the configuration of devices running
in Apstra blueprints. Before proceeding, ensure you have read the
Release Notes as well as the Upgrade Paths and Workflow documentation at:
https://www.juniper.net/documentation/product/us/en/apstra/#cat=release\_notes
https://www.juniper.net/documentation/product/us/en/apstra/#cat=install/upgrade\_software/

BLUEPRINT: 3db44826-807f-4ab9-8ca0-e25040af7ef6
(BP2)

BLUEPRINT: 964211f7-7f3c-4b0a-b6b7-137790c461f5
(BP1)

Section: FULL_CONFIG
=====
Full configuration apply.

Configuration      Role          Systems
=====
Spine      spine2 [525400E3EF4A, 10.28.105.10]
          spine1 [52540006D434, 10.28.105.9]
```

```
-----
Leaf      12-virtual-ext-001-leaf1 [5254006260B2, 10.28.105.11]
          12-virtual-ext-002-leaf1 [5254009D09D6, 10.28.105.12]
```

Warnings: Template '\_L2 Virtual EVPN' (id: '7bc432ed-c219-4e77-b08d-889ccf939add') has external connectivity settings in RackTypes: ('L2 Virtual Ext' (id: 'L2\_Virtual\_External')) which will be removed during upgrade.

The **Apstra Upgrade Summary** shows information separated by device roles (superspine, spine, leaf, leaf pair, and access switch for example). If an incremental config was applied instead of a full config, more details are displayed about the changes.

4. After you've reviewed the summary, enter **q** to exit the summary. The **AOS Upgrade: Interactive Menu** appears where you can review the exact configuration change on each device. If you're using configlets, verify that the new configuration pushed by the upgrade does not conflict with any existing configlets.



**CAUTION:** The Apstra Reference Design in the new Apstra release may have changed in a way that invalidates configlets. To avoid unexpected outcomes, verify that your configlets don't conflict with the newly rendered config. If you need to update your configlets, quit the upgrade, update your configlets, then run the upgrade again.

AOS Upgrade: Interactive Menu

```
=====
<Device SN> - display config changes using a
               specific device serial number
(s)ummary - display config change summary
(l)ist     - list all devices with config changes
(d)ump     - dump all config changes to a file
(c)ontinue - continue with AOS upgrade
(q)uit     - quit AOS upgrade
```

aos-upgrade (h for help)#

5. If you want to continue with the upgrade after reviewing pending changes, enter **c**.
6. If you want to stop the upgrade, enter **q** to abort the process. If you quit at this point and later decide to upgrade, you must start the process from the beginning.



**NOTE:** If the Apstra upgrade fails (or in the case of some other malfunction) you can gracefully shut down the new Apstra server and re-start the old Apstra server to continue operations.

## Step 4: Import State for Intent-Based Analytics



**NOTE:** As of Apstra 5.0.0, we no longer assign tags to probes and stages or support the evpn-host-flap-count telemetry service.

To remove or disable any widgets or probes for intent-based analytics, add the following arguments to the `sudo aos_import_state` command.

- To remove widgets that are not being used in any dashboards, add the `--iba-remove-unused widgets` argument.
- As of Apstra 5.0.0, we no longer assign tags to probes and stages. To remove tags from probes and stages, add the `--iba-remove-probe-and-stage-tags` argument.
- To add a serial number to non-unique probe labels, add the `--iba-number-non-unique-probe-labels` argument.
- To add a serial number to non-unique dashboard labels, add the `--iba-number-non-unique-dashboard-labels` argument.
- As of Apstra 5.0.0, we no longer support the evpn-host-flap-count service. To disable any non-predefined probes that use this service, add the `--iba-disable-probe-with-evpn-host-flap-count-service` argument.
- To remove any leading and/or trailing white space(s) from dashboard labels and their widget labels, add the `--iba-strip-dashboard-labels-widget-labels` argument.
- To remove any leading and/or trailing white space(s) from probe labels and their processor names, add the `--iba-strip-probe-labels-processor-names` argument.

## Step 5: Keep Old VM's IP Address (Optional)

If you want to keep the old VM's IP address you must perform the following extra steps before changing the Operation Mode and upgrading the devices' agent.

1. Shutdown the old VM or change its IP address to a different address to release the IP address. This is required to avoid any duplicated IP address issue.

2. Go to the new VM's Apstra interactive menu from the CLI.

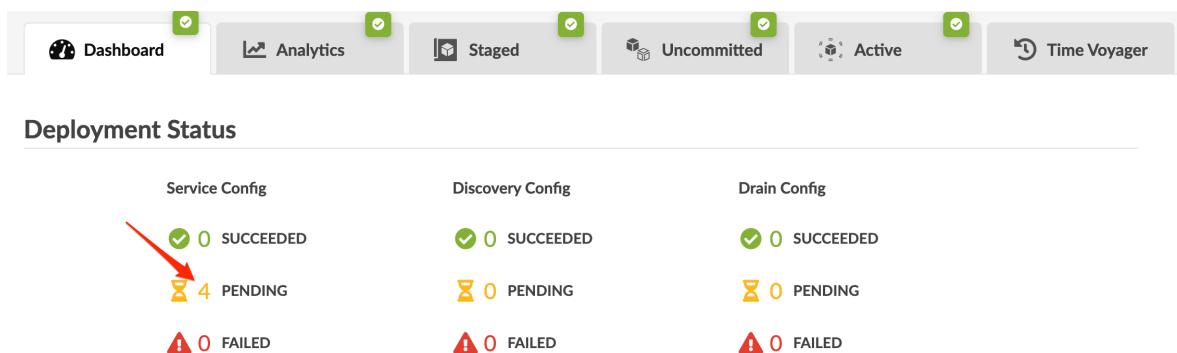
```
admin@aos-server:~$ sudo aos_config
```

3. Click **Network** to update the IP address and confirm the other parameters.
4. For the new IP address to take effect, restart the network service, either from the same menu before exiting or from the CLI after leaving the menu.

## Step 6: Change Operation Mode to Normal

When you initiate an Apstra server upgrade, the operation mode changes from **Normal** to **Maintenance** automatically. Maintenance mode prevents any offbox agents from going online prematurely. No configuration is pushed and no telemetry is pulled. At this point, if you decide to continue using the previous Apstra version instead of upgrading, you could just shut down the new Apstra server. If you decide to complete the upgrade, change the mode back to **Normal**.

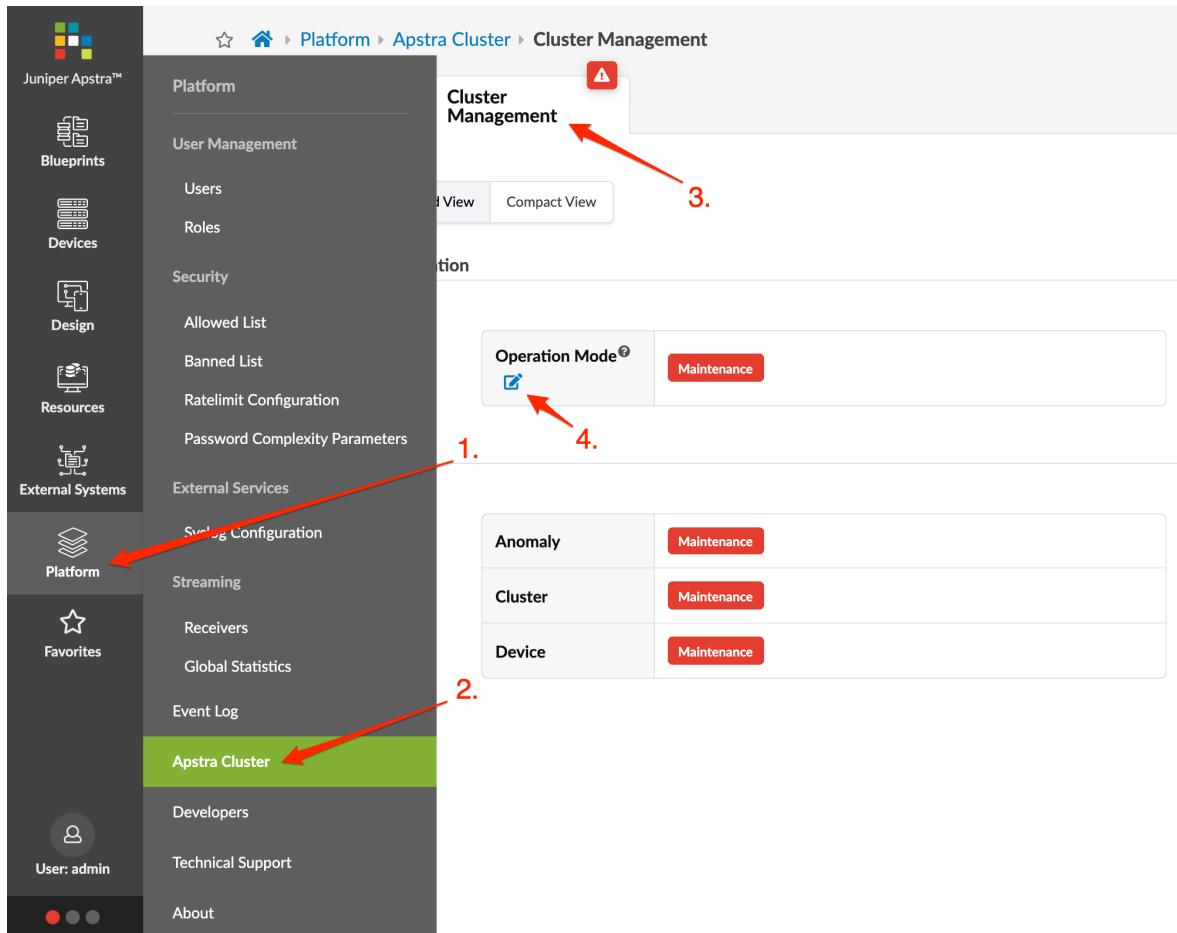
1. Log in to the Apstra GUI.
2. If you'd like to view pending service configuration changes, navigate to the dashboard of the blueprint and click **PENDING** to see the affected devices.



The screenshot shows the Apstra GUI dashboard with the following layout:

- Top Navigation:** A row of buttons: Dashboard (selected), Analytics, Staged, Uncommitted, Active, and Time Voyager.
- Section Header:** Deployment Status
- Service Config:**
  - SUCCEEDED:** 0 (green checkmark icon)
  - PENDING:** 4 (yellow warning icon with a red arrow pointing to it)
  - FAILED:** 0 (red exclamation mark icon)
- Discovery Config:**
  - SUCCEEDED:** 0 (green checkmark icon)
  - PENDING:** 0 (yellow warning icon)
  - FAILED:** 0 (red exclamation mark icon)
- Drain Config:**
  - SUCCEEDED:** 0 (green checkmark icon)
  - PENDING:** 0 (yellow warning icon)
  - FAILED:** 0 (red exclamation mark icon)

3. From the left navigation menu, navigate to **Platform > Apstra Cluster > Cluster Management**.



4. Click the **Change Operation Mode** button, select **Normal**, then click **Update**. Any offbox agents, whether they're on the controller or worker VMs automatically go online and reconnect devices and push any pending configuration changes. After a few moments the temporary anomalies on the dashboard resolve and the service configuration section shows that the operation has **SUCCEEDED**.

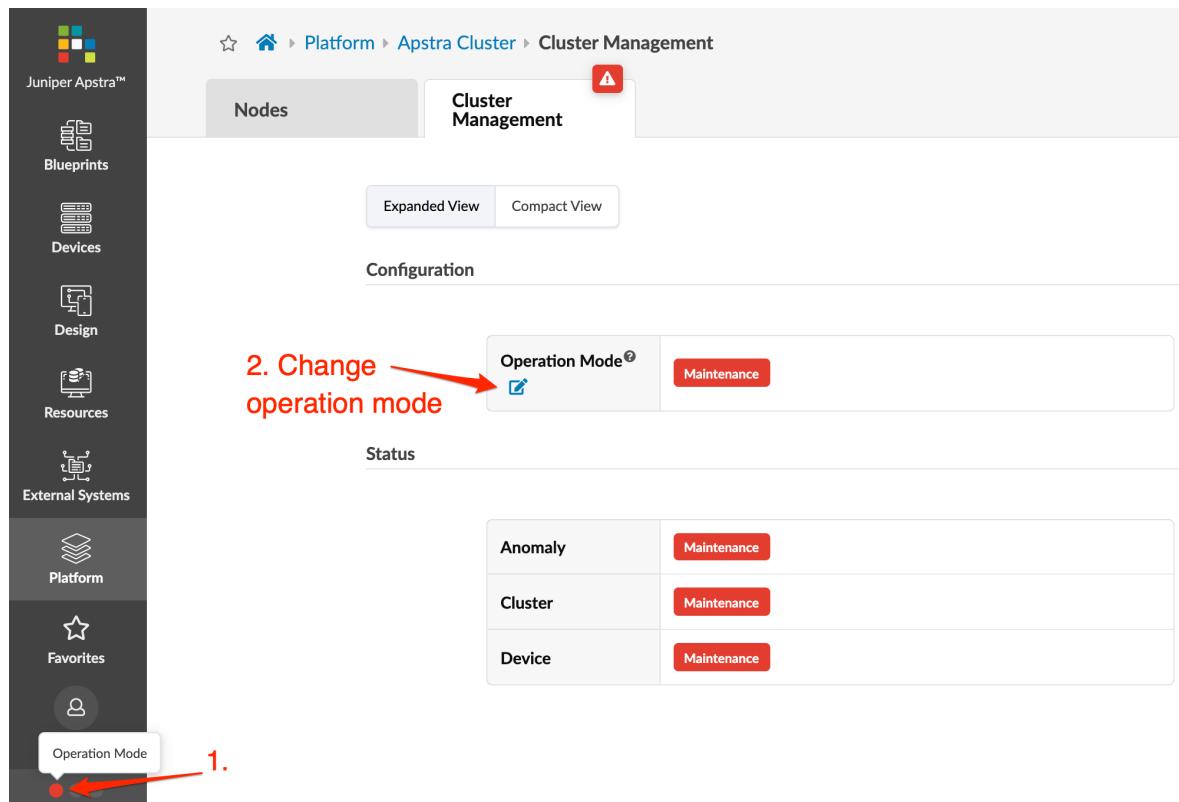


#### Deployment Status

Service Config	Discovery Config	Drain Config
<span>✓</span> 4 SUCCEEDED	<span>✓</span> 0 SUCCEEDED	<span>✓</span> 0 SUCCEEDED
<span>✗</span> 0 PENDING	<span>✗</span> 0 PENDING	<span>✗</span> 0 PENDING
<span>⚠</span> 0 FAILED	<span>⚠</span> 0 FAILED	<span>⚠</span> 0 FAILED

You can also access the **Cluster Management** page from the lower left section of any page. You have continuous platform health visibility from here as well, based on colors.

From the bottom of the left navigation menu, click one of the dots, then click **Operation Mode** to go to **Cluster Management**. Click the **Change Operation Mode** button, select **Normal**, then click **Update**.



## Step 7: Upgrade Onbox Agents

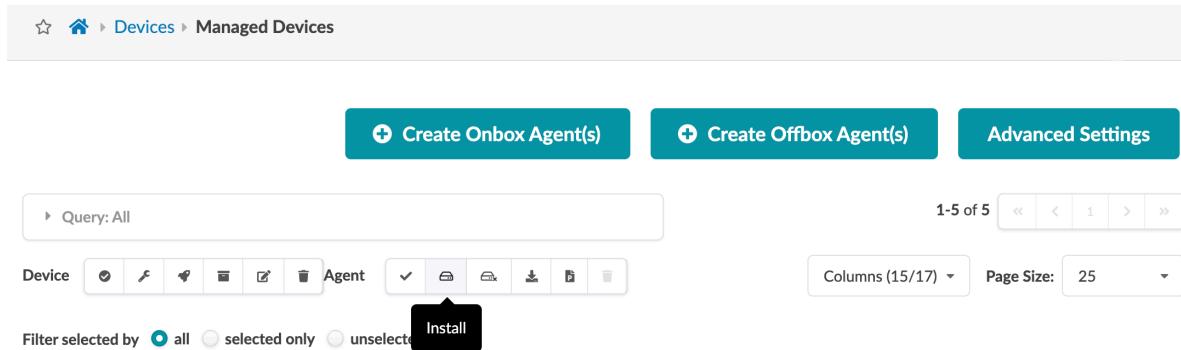
The Apstra server and onbox agents must be running the same Apstra version. If versions are different the agents won't connect to the Apstra server.

If you're running a multi-state blueprint, especially 5-stage, we recommend that you upgrade agents in stages: first upgrade superspines, then spines, then leafs. We recommend this order because of path hunting. Instead of routing everything up to a spine, or from a spine to a superspine, it's possible for routing to temporarily go from leaf to spine back down to another leaf and back up to another spine. To minimize the chances of this happening, we recommend upgrading devices in stages.

1. Log in to the Apstra GUI as user **admin**.
2. From the left navigation menu, navigate to **Devices > Managed Devices** and select the check boxes for the device(s) to upgrade (up to 100 devices at a time). You can upgrade multiple onbox agents at the same time, but the order of device upgrade is important.
  - Upgrade agents for superspines first.
  - Upgrade agents for spines second.
  - Upgrade agents for leafs third.

When you select one or more devices the **Device** and **Agent** menus appear above the table.

3. Click the **Install** button to initiate the install process.



The screenshot shows the Juniper Apstra interface with the 'Managed Devices' page selected. At the top, there are three buttons: 'Create Onbox Agent(s)', 'Create Offbox Agent(s)', and 'Advanced Settings'. Below the buttons is a search bar with the placeholder 'Query: All'. To the right of the search bar is a page navigation area showing '1-5 of 5' with arrows for navigation. Underneath the search bar is a toolbar with various icons for device management, including a 'Device' dropdown and several small icons. Below the toolbar is a filter section with radio buttons for 'all', 'selected only', and 'unselected'. The 'Install' button is highlighted with a black box and an arrow pointing to it. The bottom right corner of the interface shows 'Columns (15/17)' and 'Page Size: 25'.

The job state changes to **IN PROGRESS**. If agents are using a previous version of the Apstra software, they are automatically upgraded to the new version. Then they connect to the server and push any pending configuration changes to the devices. Telemetry also resumes, and the job states change to **SUCCESS**.

4. In the **Liveness** section of the blueprint dashboard confirm there are no device anomalies.



**NOTE:** If you need to roll back to the previous Apstra version after initiating agent upgrade, you must build a new VM with the previous Apstra version and restore the configuration to that VM. For assistance, contact Juniper Technical Support.

## Step 8: Shut Down Old Apstra Server

1. Update any DNS entries to use the new Apstra server IP/FQDN based on your configuration.
2. If you're using a proxy for the Apstra server, make sure it points to the new Apstra server.
3. Gracefully shut down the old Apstra server. You will have been asked if you want the old Apstra server shut down; if you responded yes, then the `aos stop` command is run automatically to shut down the old Apstra server for you.
4. If you're upgrading an Apstra cluster and you replaced your worker nodes with new VMs, shut down the old worker VMs as well.

Next Steps:

If the NOS versions of your devices are not qualified on the new Apstra version, upgrade them to a qualified version. (See the [Juniper Apstra User Guide](#) for details.)

## Roll Back Apstra Server Upgrade

If you've upgraded the Apstra server onto a different VM from the previous version, you can roll back to the previous version. (If you've upgraded on the same virtual machine, this option is not available.) You'll lose any changes that you've made on the new Apstra server since upgrading. This action is disruptive.



**CAUTION:** Apstra does NOT support 'hitless' rollbacks.

1. Shut down the VM containing the 'new' Apstra version.
2. Install / downgrade device onbox agents to match the 'old' Apstra version. (Offbox agents will take care of themselves.)
3. Apply full configuration to the network devices to restore configuration from the 'old' Apstra server.

The screenshot shows the Apstra interface with the 'Active' tab selected. Under the 'Physical' tab, the 'Config' section is highlighted. A red arrow points to the 'Apply Full Config' button, which is highlighted with a red box. The status bar at the bottom indicates 'Actual config deviated from golden config'.

Intended running configuration		Actual running configuration
1   Command: show running-config	1   Command: show running-config	1   Command: show running-config
2   device: IEDBL4-01BLF01 (DCS-7280SR2-48YC6, EOS-6.24.3M)	2   device: IEDBL4-01BLF01 (DCS-7280SR2-48YC6, EOS-6.24.3M)	2   device: IEDBL4-01BLF01 (DCS-7280SR2-48YC6, EOS-6.24.3M)
3	3	3