



nHCI Installation Guide

Clustering with HA and virtualized nSAN storage

Contents

1	Overview	2
2	Install nHCI and Perform Initialization	2
2.1	Hardware Requirements	2
2.2	Create Bootable USB Drive	4
2.3	Installation	4
2.4	Installation Summary Page	4
2.5	Select Installation Mode	5
2.6	Configure Disk Partitions	6
2.7	Configure Network	6
2.8	Root Password Configuration and Installation Start	14
2.9	Begin The Installation Process	14
3	Cluster Initialization	14
3.1	Create Data Center	15
3.2	Create Cluster	15
3.3	Add Host to Cluster	16
4	High Availability Configuration	17
4.1	Accessing the Management Node Operations	17
4.2	Adding the Second Management Node	18
4.3	Accessing the Cluster via the Virtual IP	20
5	Distributed Storage (nSAN) Installation	21
5.1	Creating the Distributed Switch	21
5.2	Distributed Port Group and Kernel Adapters	24
5.3	Uploading the nSAN Package	25
5.4	Adding Storage Nodes to nSAN	29
5.5	Initializing Physical Disks and Creating Data Disks	31
5.6	Creating the Storage Pool	33
5.7	Registering the Storage Pool and Creating the Cluster Primary Storage	35
6	Final Cluster Preparations	38
6.1	Configure Image Storage	38
6.2	Configure Network Resources	40
7	Creating Virtual Machines	41
8	Conclusion of the Installation	41

1 Overview

This document describes how to install **nHCI**.

This installation guide focuses on **practical implementation** rather than serving as an exhaustive technical reference. The objective is to provide administrators and engineers with clear, step-by-step instructions that can be followed in real deployment scenarios.

The guide covers the most common **clustering configurations**, including setups with **High Availability (HA)** and **virtualized storage with nSAN**. It assumes the use of **standard hardware resources** and typical network environments, so that the procedures described can be reproduced in most datacenter or lab contexts without requiring custom adaptations.

Advanced tuning, troubleshooting, and edge-case scenarios are outside the scope of this document.

The guide provides:

- Step-by-step installation procedures for nHCI
- Essential hardware and network configuration requirements
- Basic post-installation setup through both CLI and web interface

Note

A functional nHCI cluster requires either a setup of two nodes plus a witness, or a configuration of three or more nodes. The specific differences between these deployment options will be detailed in the following sections of this guide.

2 Install nHCI and Perform Initialization

2.1 Hardware Requirements

The configuration of server CPU, memory, storage size, and NIC performance determines the business capacity of the nSSV.

Production Environment:

- **Management Node Configuration:** Determined by deployment scale and user environment, please consult official technical support for details. For small-scale scenarios, the recommended configuration for the management node is: 8 CPU cores, 16 GB of memory, and 240 GB of storage.
- **Compute Node Configuration:** Determined by business scale, please consult official technical support for details.

Recommended Configurations for Server Hardware Whether for minimum environment deployment or production environment deployment, it is recommended that hardware devices such as servers be configured according to the recommended configurations in the table below:

Device	Component	Configuration Requirements
Server	CPU	x86 Environment: 64-bit x86 architecture, supporting Intel VT or AMD-V hardware virtualization features (for example, Intel's VMX or AMD/Hygon's SVM). ARM Environment: 64-bit ARM architecture, supporting hardware virtualization features.
	Memory	No special requirements. Recommends DDR4 or higher performance memory.
	Motherboard	Standard dual-socket server motherboard.
	RAID Card	Supports SAS/SATA RAID 0/1/10 and supports passthrough mode.
	Hard Disk	No special requirements. You can choose between HDD or SSD based on storage planning.
	Network Port	<ul style="list-style-type: none"> • 1 Gigabit Ethernet port for management network, for example, Ethernet 1GbE, RJ45 • 1 10 Gigabit Ethernet port for business network, for example, Ethernet 10GbE, SFP+
Network Switch	–	<ul style="list-style-type: none"> • At least 1 Gigabit switch, 10 Gigabit switch recommended • Several Category 5 cables

Table 1: System Configuration Requirements

Note

The nHCI operating system must be installed on solid-state storage.

When preparing your server hardware, make sure to configure the following:

- Enable CPU virtualization support in the server BIOS.
- Complete your storage planning in advance:
 - If you use local storage, it is recommended to adopt a storage redundancy backup solution (for example, configure 4 hard disks in RAID 10) to improve the reliability of data storage and image storage. If your virtual machines require very high I/O read/write performance, consider using a RAID configuration with all SSDs. If the I/O access of your virtual machines leans more towards read performance, a mixed configuration of SSDs and HDDs can also work well.
 - If you use NFS or distributed storage, configure the corresponding storage or file system in advance. If your image storage uses a distributed image storage, ensure that your data storage also uses distributed storage.
- Plan your network in advance: It is recommended that you consistently name all physical host NICs and use NICs with the same name to carry the same type of communication traffic. For example, management traffic should all use the em1 NICs.
- Complete the necessary configurations on your network switch in advance:

- If you need to use a VLAN network environment, configure the corresponding VLAN network communication on the switch in advance.
- nSSV will actively allocate IP addresses to your virtual machines, so reserve a range of IP addresses that do not conflict with your system, and avoid conflicts with existing DHCP services in your network environment.

2.2 Create Bootable USB Drive

To begin the installation, you need to create a bootable USB drive with the nSSV ISO image:

- Download the nSSV ISO image from the official repository
- Use a tool like Rufus to write the ISO to a USB drive
- Ensure the USB drive has at least 8GB of capacity
- Insert the bootable USB drive into the target server
- Configure the server BIOS/UEFI to boot from the USB device

The system will boot from the USB drive and start the nHCI installation process.

2.3 Installation

Enter the ISO boot interface and choose the default option to start the operating system installation. You can select based on your actual situation, but we recommend using the graphical user interface (GUI) for installation. If the server does not have a VGA port and only supports serial connections, you can use either VNC or text mode installation methods.

2.4 Installation Summary Page

This page displays the system installation configuration. You can modify the configuration as needed.

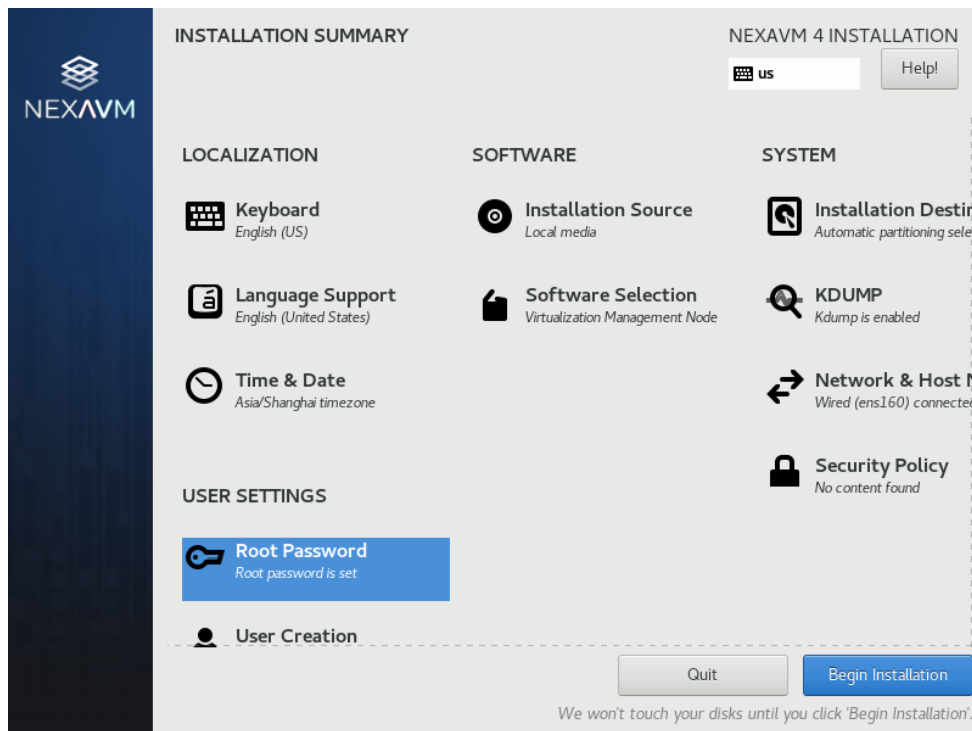


Figure 1: Installation Summary

2.5 Select Installation Mode

On the *Installation Summary* page, click *Software Selection*.

On the *Software Selection* page, choose the installation mode based on the intended role of the node within the cluster:

- *Management Node* – select this option if the current server will act as a **management node**. This node will host the nSSV management services and, in High Availability (HA) deployments, will synchronize with a peer management node.
- *Compute Node* – select this option if the current server will serve as a **compute node** (or **witness node**) within the cluster. Compute nodes provide virtualization and processing resources, while witness nodes are required in two-node HA configurations to ensure proper arbitration and quorum for nSAN storage.

For the purpose of this guide, make sure to select **Management Mode** for at least two nodes. Since an nHCI cluster requires a high-availability management plane, a minimum of two management nodes must be deployed. The remaining nodes in the cluster can be configured as compute nodes or, where applicable, as witness nodes.

Note

In HA setups with virtualized storage (**nSAN**), at least two management nodes plus one witness are required. Select *NexaVM Compute Node* for the witness node installation.

After selecting the appropriate installation mode, click **Done** to return to the Installation Summary page.

2.6 Configure Disk Partitions

On the *Installation Summary* page, click *Installation Destination* to enter the *Installation Destination* page.

Note

We recommend that you only configure the system disk on the page. After the system is installed, you can configure other disks.

For Device Selection, we recommend that you only configure the system disk. After the system is installed, you can configure other disks.

If the selected disk does not have enough available space, click *Reclaim Space* and *Delete All*.

For *Storage Configuration*, we recommend selecting *Automatic* to automatically configure the disk partitions.

If you need to manually configure disk partitions, refer to the following guidelines based on the BIOS boot mode:

UEFI Mode:

- `/boot`: This directory stores the core files needed for Linux boot. We recommend allocating 1GB of space.
- `/boot/efi`: This directory stores the UEFI boot files. We recommend allocating 500MB.
- `/`: This is the root directory for the Linux system. We recommend allocating all remaining space.

Legacy Mode:

- `/boot`: This directory stores the core files needed for Linux boot. We recommend allocating 1GB of space.
- `/`: This is the root directory for the Linux system. We recommend allocating all remaining space.

Note

- The above values represent the recommended partition sizes for nHCI (total disk capacity should be greater than 300GB).
- In Legacy mode, if the system disk capacity exceeds 2TB, you need to configure a BIOS boot partition to support GPT partitioning. UEFI mode does not have this limitation and supports GPT partitioning.

Review the configuration and click *Done*.

2.7 Configure Network

Bond NIC Deployment

To meet the bandwidth capacity and high reliability requirements, NIC bonding must be implemented in the production environment. The operating system installation process is identical

for both x86 and ARM servers. The following sections describe the installation steps using an x86 server as a reference.

The procedure shown below illustrates, for simplicity, the creation of the management bond using a single NIC, as shown in the example screenshots. This approach is intended to clearly demonstrate the fundamental steps without introducing additional complexity related to multi-interface configurations.

The remaining network connections (for example, those dedicated to storage or other types of traffic) will be configured later through the nHCI GUI, according to the system architecture and performance requirements.

Note

Pay attention to disabling all IPv4 addresses on the various NICs, as will be shown later in the screenshots.

Add a bond device.

- a) On the **INSTALLATION SUMMARY** page, click **Network & Host Name**.
- b) On the **NETWORK & HOST NAME** page, click the **+** button at the bottom left of the page. This will open the *Add device* dialog. From the drop-down list, choose **Bond**, then click **Add**.

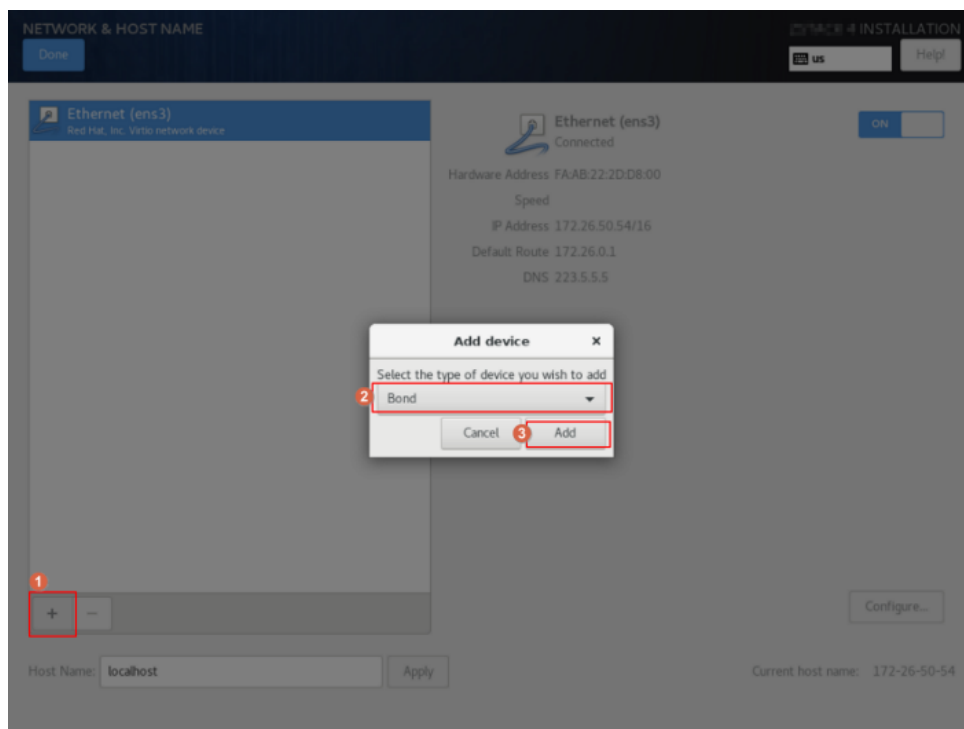


Figure 2: Add Bond Device

Add a Bond Slave.

- a) In the Bond configuration dialog, adjust the **Connection name** as needed.

Note

Connection name and interface name **must** match as showed in the screenshot below.

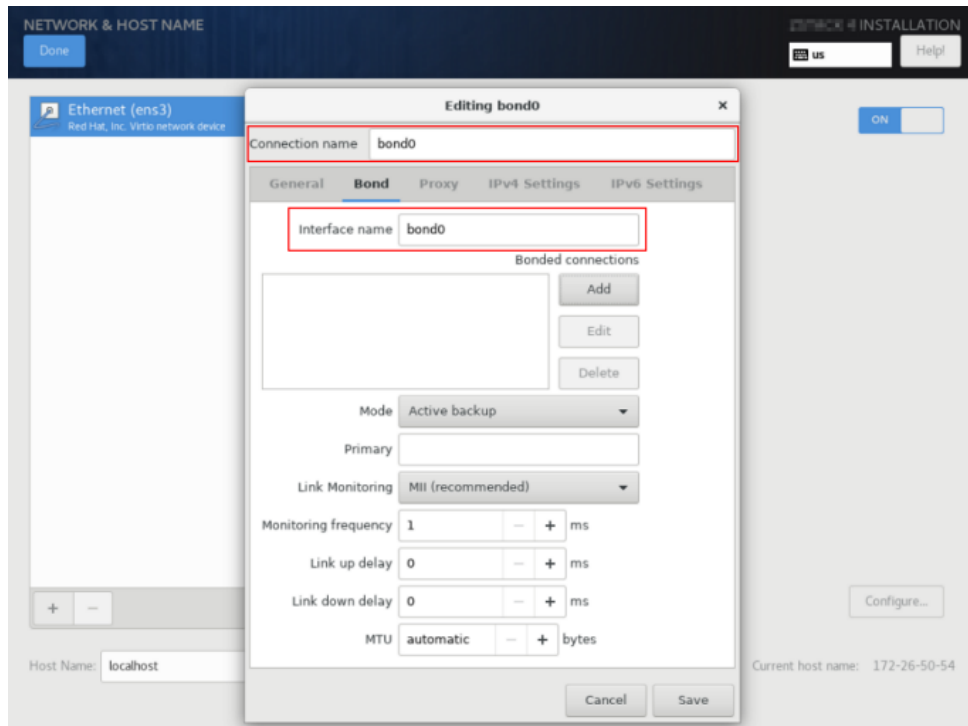


Figure 3: Adjust Connection Name

b) On the Bond configuration window, click **Add** to add a bond slave.

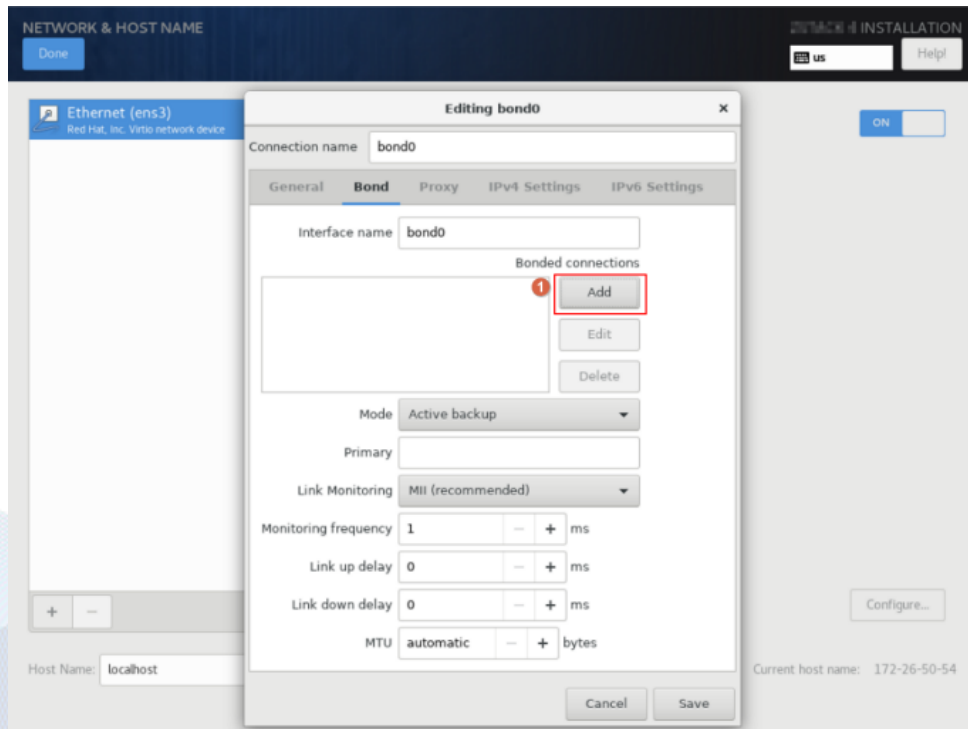


Figure 4: Add Bond Slave

- c) In the *Choose a Connection Type* dialog, choose a connection type from the drop-down list, such as **Ethernet**, and then click **Create**.

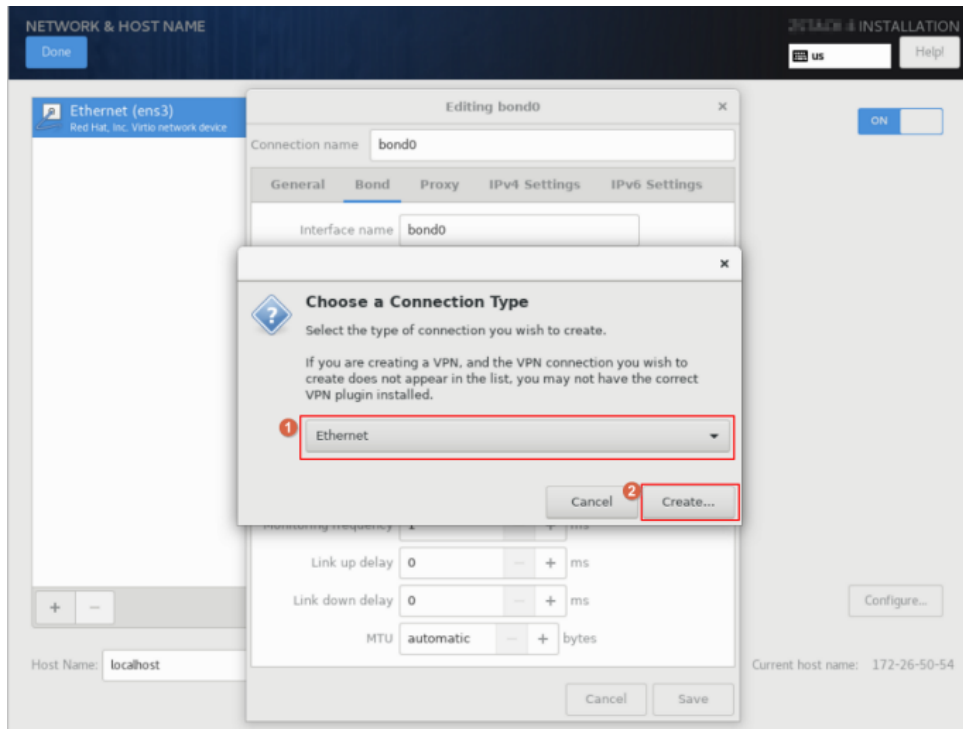


Figure 5: Select Bond Slave Connection Type

- d) On the **Ethernet** tab of the *Editing bond0 slave1* dialog, click **Device** and select the Slave device you want to bind, such as **ens3** (corresponding MAC address). Keep the other options as default or customize them as needed, then click **Save**.

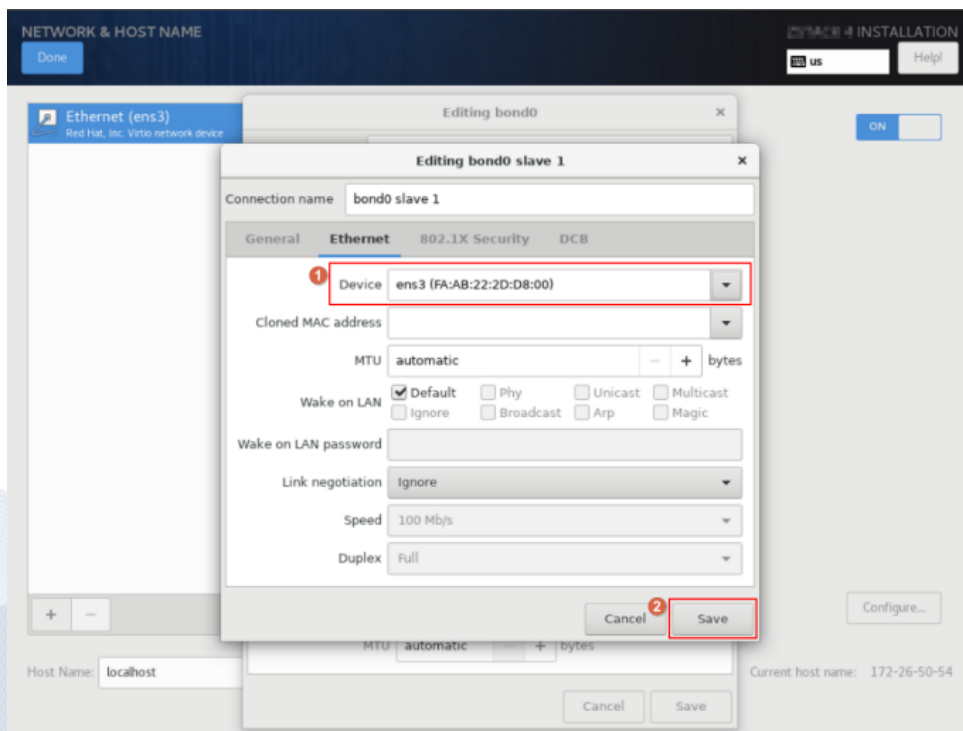


Figure 6: Select Bond Slave Device

Select the Bond mode.

In the Bond configuration dialog, choose the bond mode from the **Mode** drop-down list as needed, such as **Active backup**. Keep the other options as default or customize them as needed, then click **Save**.

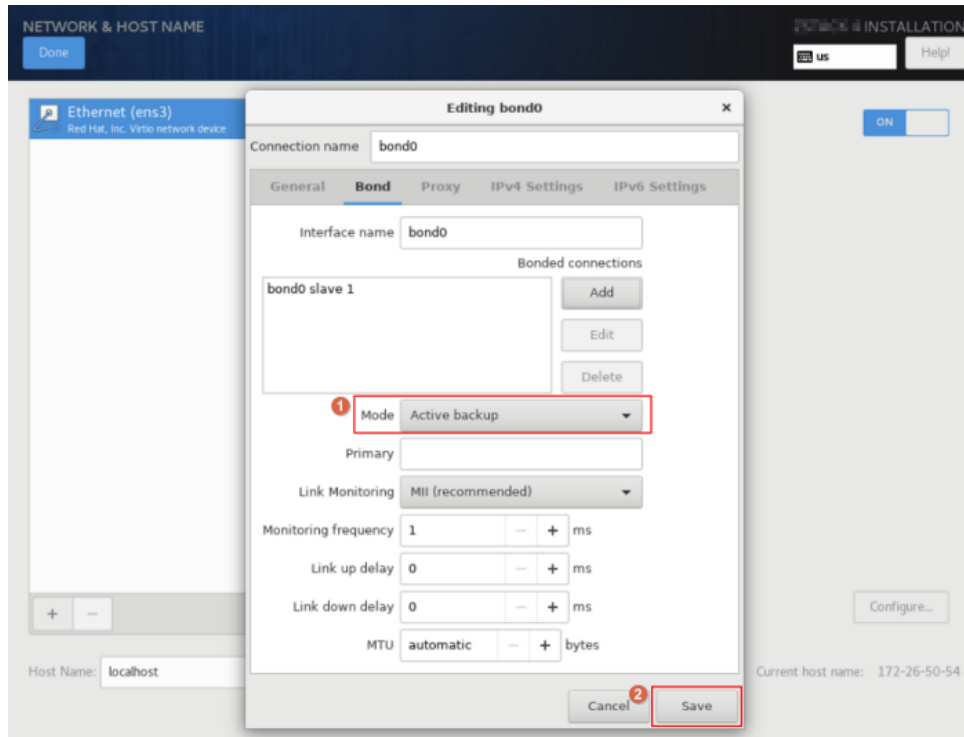


Figure 7: Select Bond Mode

Disable IPv4 on the original NIC.

- a) Select the original NIC, such as **Ethernet (ens3)**, and click **Configure**.

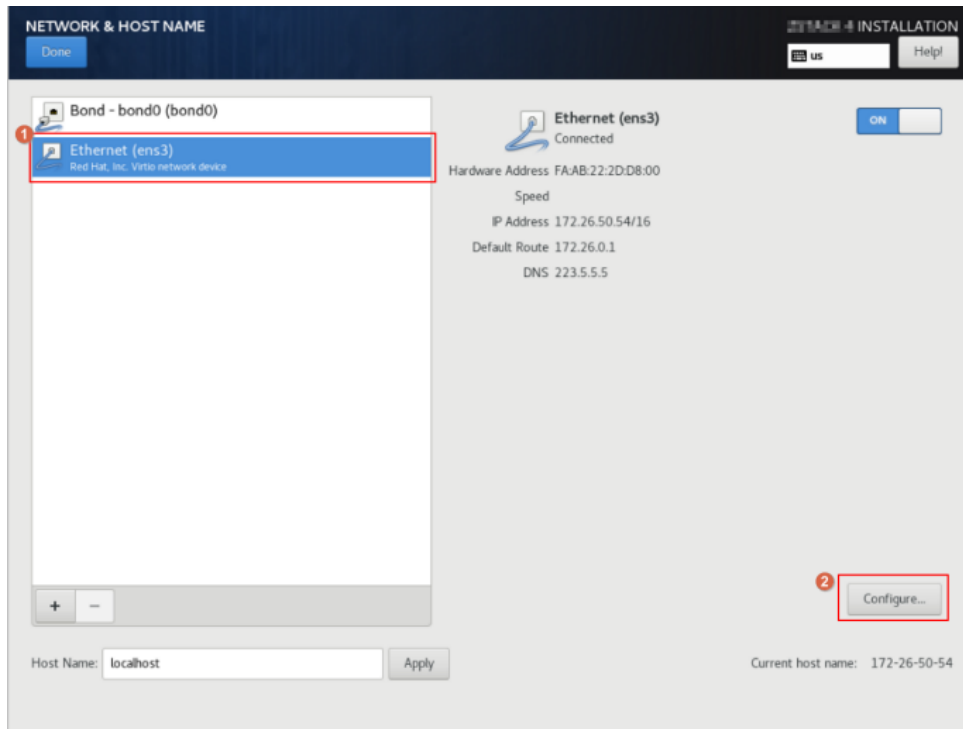


Figure 8: Configure Original NIC

- b) The *Editing ens3* dialog pops up. Click **IPv4 Settings** to access the IPv4 Settings tab. Change the **Method** parameter value to **Disabled**, then click **Save**.

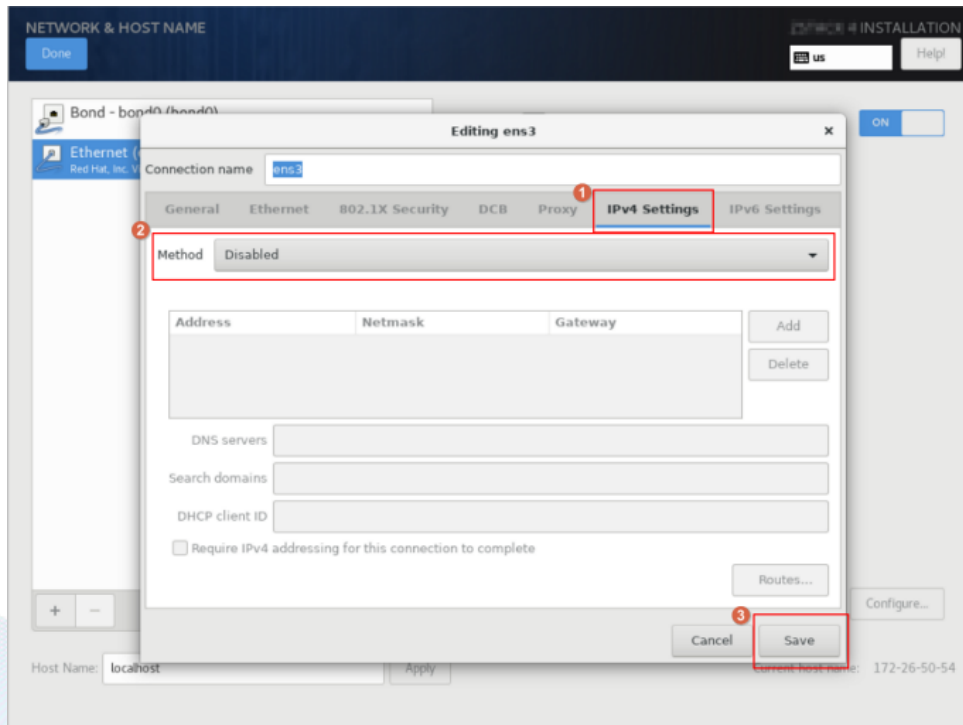


Figure 9: Disable IPv4

Configure a static IP address for Bond.

- a) On the **NETWORK & HOST NAME** page, choose the bond device, such as **Bond (bond0)**, and then click **Configure**.
- b) The *Editing bond0* dialog pops up. Click **IPv4 Settings** to access the IPv4 Settings tab. Change the **Method** parameter value to **Manual** to switch the IP address acquisition method to manual.

Note

You can configure the IP address acquisition method as needed, including using DHCP for automatic acquisition or specifying it manually.

- c) Click **Add** to add an IP address entry. Configure the IP address, netmask, and gateway as needed, then click **Save** to save the configuration.

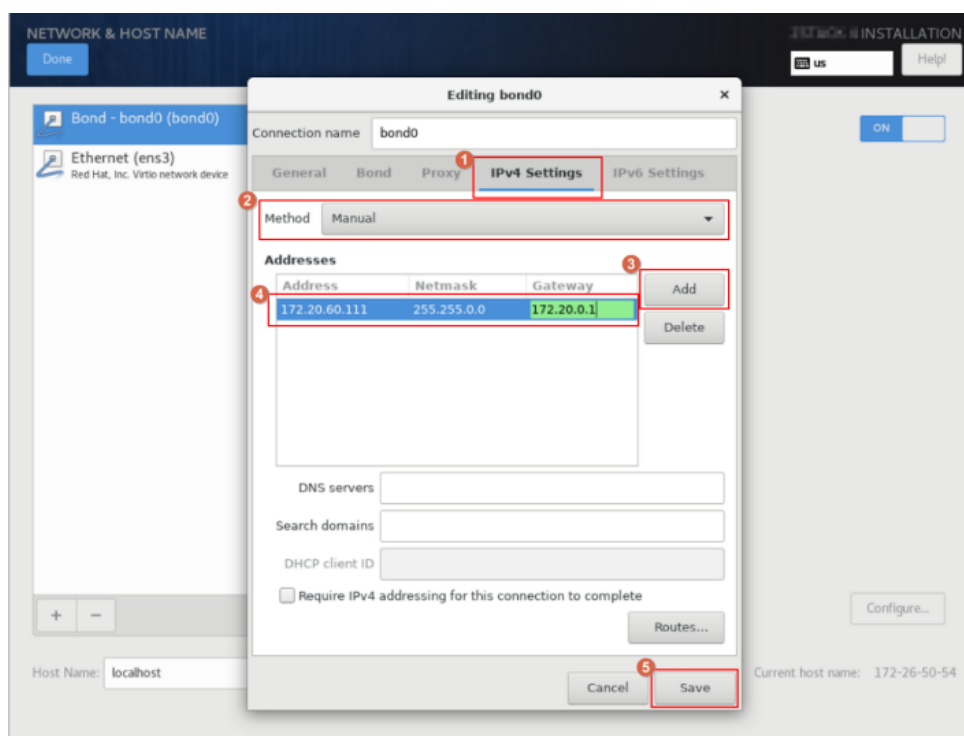


Figure 10: Configure Bond Static IP Address

Configure the NIC to activate automatically.

In the *Editing bond0* dialog, click **General** to access the General tab. Select the **Connect automatically with priority** checkbox to set the NIC to activate automatically, then click **Save**.




Figure 11: Configure NIC to Activate Automatically

Complete the Bond configuration.

- Check the Bond configuration settings. Ensure that **On** is enabled and that you have configured the IP address. Also, make sure the Bond Slave (such as **ens3**) is set to **On**. Otherwise, nSSV will not be installed properly.
- Review the configuration and click **Done** to return to the **INSTALLATION SUMMARY** page.

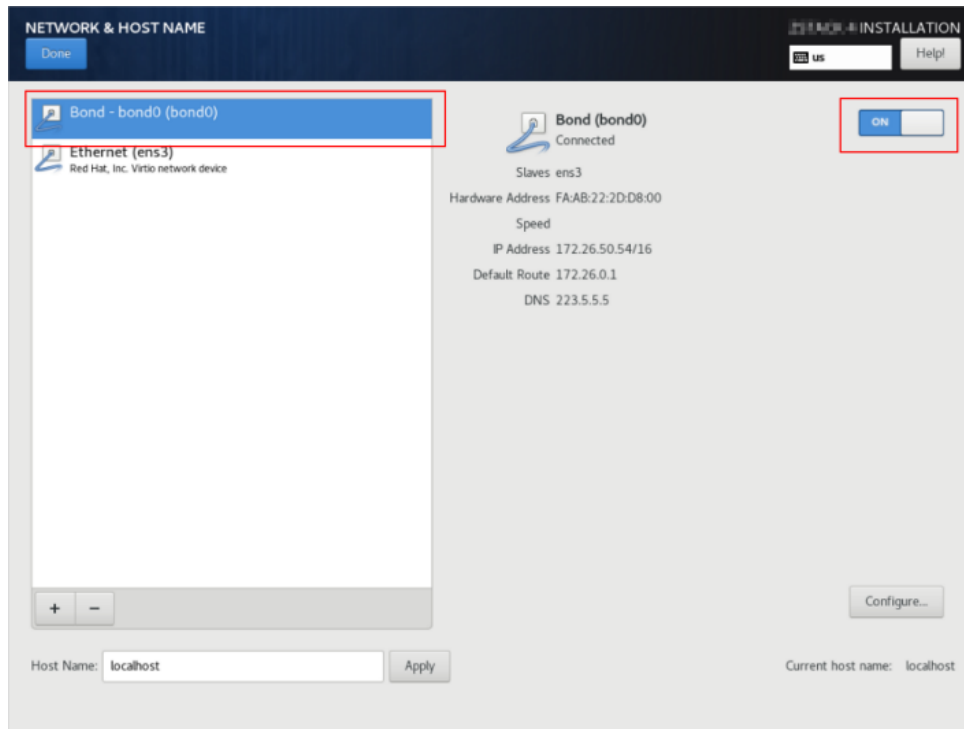


Figure 12: Check Bond Configuration

2.8 Root Password Configuration and Installation Start

From the **INSTALLATION SUMMARY** page, configure the root account by selecting **Root Password** and entering the desired password for the operating system. Ensure that the password complies with the system security requirements.

Once the root password has been configured, return to the **INSTALLATION SUMMARY** page and select **Begin Installation** to start the operating system installation process.

2.9 Begin The Installation Process

Once all required configurations have been completed, click **Begin Installation** in the bottom-right corner of the Installation Summary page.

Wait for the first part of the installation to complete. When the screen turns black and the server begins to reboot, **remove the USB drive** to prevent the system from booting from it again.

The system will now complete the installation and reboot. After the reboot, we will continue with the backend configuration manually through the command line interface.

3 Cluster Initialization

After installing nSSV on all servers, access the web interface of one of the two management nodes by opening a browser and navigating to:

```
https://[MANAGEMENT_NODE_IP]
```

Log in using the default credentials:

- **Username:** admin
- **Password:** password

Upon the first login, the system will display the initialization wizard. In a standard workflow this tool is used to perform the first setup steps of the environment. However, in a hyperconverged deployment the virtual storage layer (nSAN) and the HA components must be installed and synchronized *before* completing the cluster initialization.

For this reason, at this stage we recommend closing the wizard by clicking **Cancel**. This allows the platform to be prepared in the correct sequence, ensuring that the subsequent cluster, host, network, and storage configurations are applied to a fully initialized HA and storage environment.

All required configuration steps will be carried out in the order described in the following sections of this guide.

3.1 Create Data Center

1. Click the **Root Node** (the IP of the in use node)
2. Click the **Actions** button and select **New Data Center**
3. Fill in the required fields
4. Click **OK** to create the data center

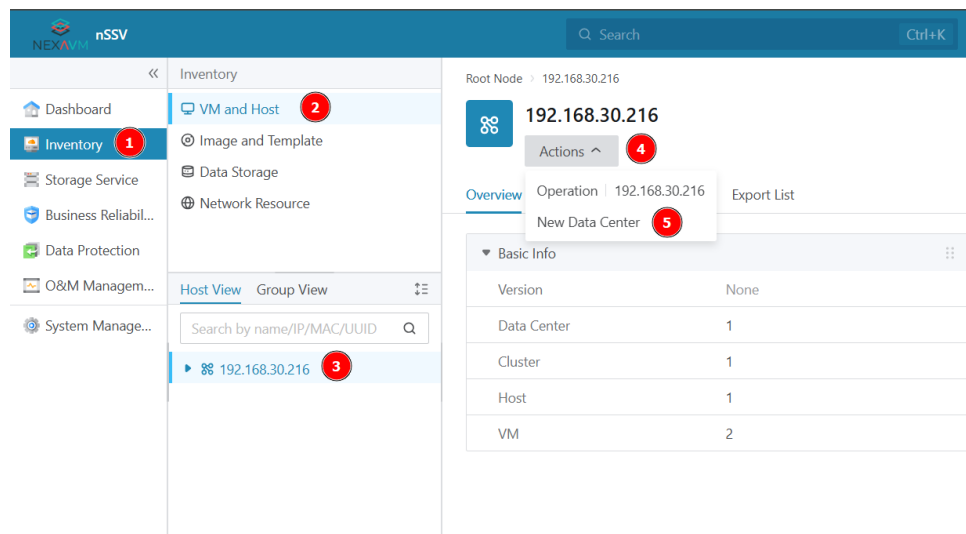


Figure 13: New Data Center Flow

3.2 Create Cluster

1. Click on the newly created **Data Center** in the left menu
2. Click the **Actions** button and select **New Cluster**
3. Fill in the required fields
4. Click **OK** to create the cluster

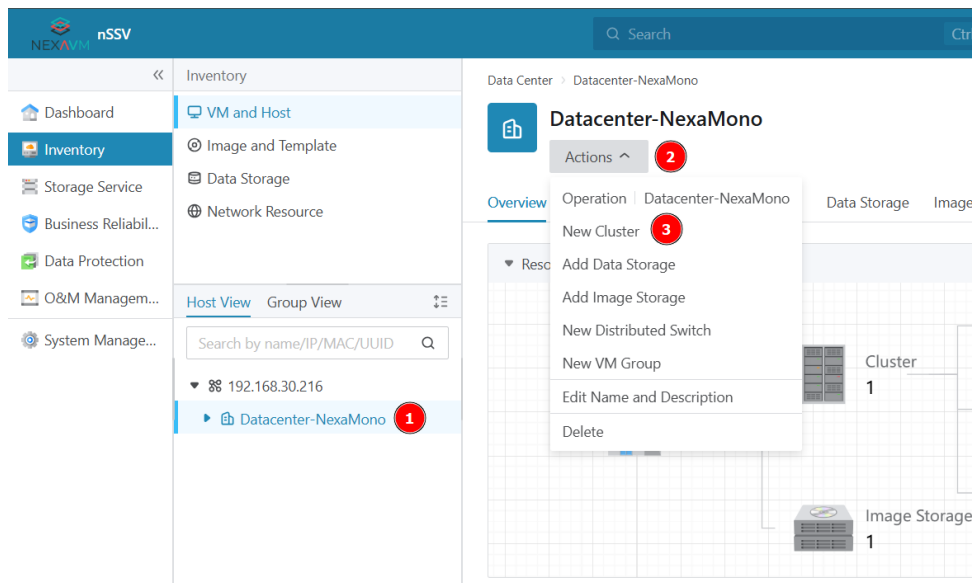


Figure 14: New Data Center Flow

3.3 Add Host to Cluster

Once the cluster has been created, the next step is to register the servers that will participate in it. The procedure described below must be performed for **each host** that will be part of the nHCI environment, including management nodes and compute nodes.

1. Click on the newly created **Cluster** in the left menu
2. Click the **Actions** button and select **Add Host**
3. Fill in the host information:
 - **Addition Method:** Select **Single**
 - **Host IP:** Enter the management IP address configured earlier
 - **SSH Port:** Enter the SSH port (default: 22)
 - **Username:** Enter the root username
 - **Password:** Enter the root password created during installation
4. Click **OK** to add the host to the cluster

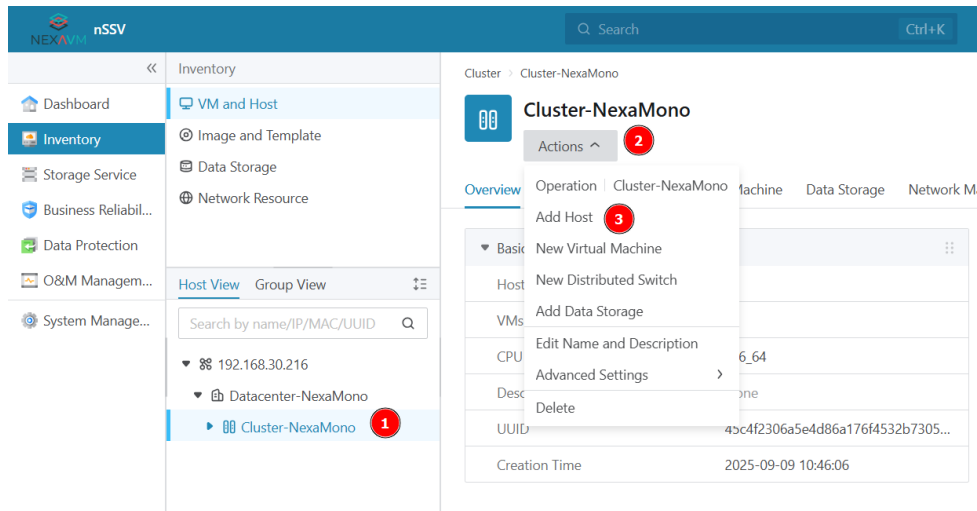


Figure 15: New Data Center Flow

4 High Availability Configuration

Once all management nodes have been added to the cluster, the next step is to enable the High Availability (HA) service for the management plane. This procedure ensures that the nHCI environment can continue operating even in the event of a failure of one of the management nodes.

4.1 Accessing the Management Node Operations

From the web interface of the management node currently in use, navigate to:

- **Reliability** → **MN Monitoring**

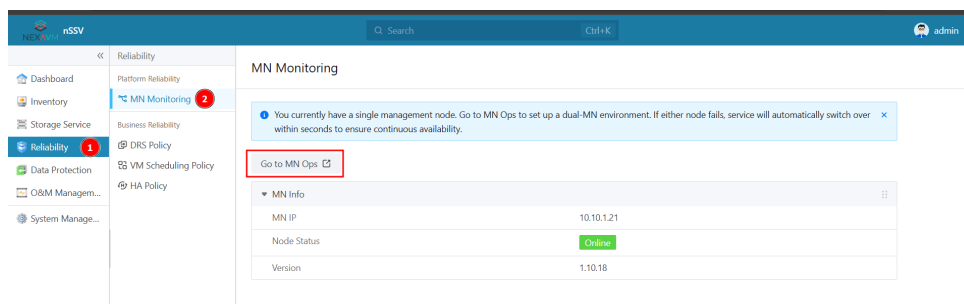


Figure 16: MN Monitoring Navigation

On the MN Monitoring page, click on **Go to MN Ops** to access the management node operations panel.

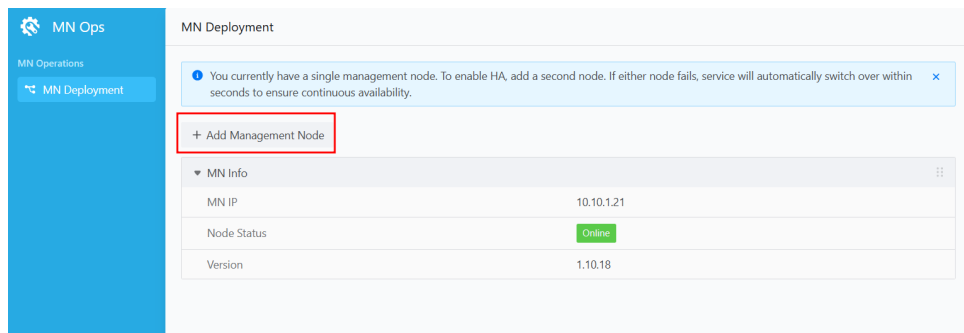


Figure 17: MN Monitoring Navigation

4.2 Adding the Second Management Node

In the MN Operations interface, click on **Add Management Node**. A configuration window will appear, prompting you to enter the information required to establish the HA relationship between the two management nodes.

Fill in the requested fields carefully:

- **VIP:** The Virtual IP that will serve as the unified entry point for management access.
- **Peer Management Node IP:** The management IP address of the second management node.
- **SSH Credentials:** Username, password, and port required to access the peer node.
- Any additional parameters requested by the wizard.

management node. To enable HA, add a second node. If either node fails, service will automatically switch over to the other node if available.

Add Management Node

Configure MN Review Configuration Add MN

VIP *	<input type="text" value="10.10.1.20"/>
Active MN	10.10.1.21
SSH Username *	<input type="text" value="root"/>
SSH Password *	<input type="password" value="....."/>
Standby MN IP *	<input type="text" value="10.10.1.22"/>
SSH Username *	<input type="text" value="root"/>
SSH Password *	<input type="password" value="....."/>
Time Svc Server	10.10.1.21

Cancel [Next >](#)

Figure 18: Add Management Node

After entering all required information, proceed through the wizard until the configuration is fully completed.

Add Management Node
×

Configure MN
Review Configuration
Add MN

Active MN		Standby MN	
Node IP	10.10.1.21	Node IP	10.10.1.22
Data Center	1	Data Center	0
Cluster	1	Cluster	0
Host	1	Host	0
VM	0	VM	0

I acknowledge the above risks. To confirm to Add, type **Add** here.

Cancel
< Back
OK >

Figure 19: Add Management Node

4.3 Accessing the Cluster via the Virtual IP

Once the HA wizard is successfully completed, a pop-up message will confirm that the management cluster is now operating under the defined Virtual IP (VIP).

At this point, access the nHCI management interface using the VIP:

```
https://[VIP_ADDRESS]
```

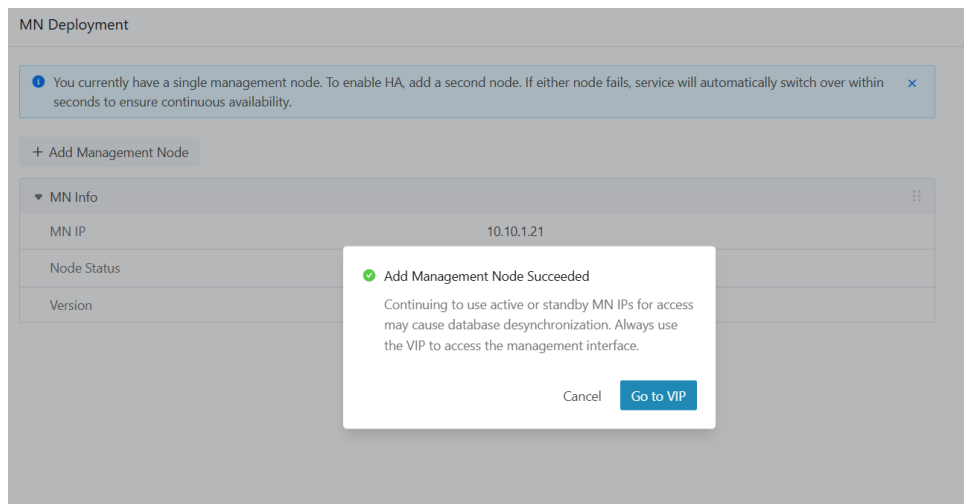


Figure 20: VIP Login

You may log in using the same credentials that were used during the first access to the nSSV web interface.

The management plane is now highly available and ready for the subsequent configuration steps of the nHCI environment.

5 Distributed Storage (nSAN) Installation

After completing the HA configuration of the management plane, the next step is to prepare the storage network and deploy the nSAN distributed storage service. This phase ensures that all hosts in the cluster can communicate through a dedicated storage network and synchronize their data volumes efficiently.

5.1 Creating the Distributed Switch

From the web interface, navigate to:

- **Inventory** → **Network Resource**

Select the previously created **Data Center**, then create a new **Distributed Switch** as shown in the reference screenshot.

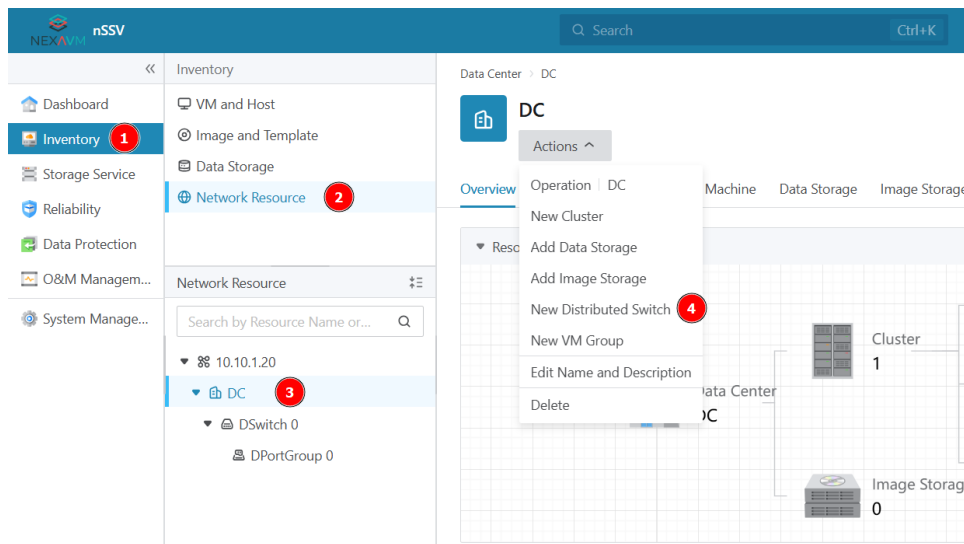


Figure 21: Distributed Switch

Follow the configuration steps presented by the wizard. During the creation process, in the **Bond** section, select the appropriate option based on your network setup:

- **Add Individually** – add network interfaces one by one to the bond.
- **Batch Bonding** – aggregate multiple interfaces simultaneously.
- **Specify Same Port** – assign a bond shared across all hosts in the cluster.

Choose the option that best suits your infrastructure requirements to ensure that all storage traffic is correctly isolated and routed through the intended interfaces.

New Distributed Switch ×

Basic Info

Name *

Description
0/256

Data Center DC

Cluster * 🗑️
+ Add Cluster

Network Configuration Distributed Port Group Configuration

Addition Method ^

- Add Individually
- Batch Bonding** ✓
- Specify the Same Port

Uplink Name *

Bond Mode

Hash Policy v

Cancel **OK**

Figure 22: Distributed Switch

New Distributed Switch ×

Basic Info

Name *

Description
0/256

Data Center DC

Cluster * 🗑️
+ Add Cluster

Network Configuration **Distributed Port Group Configuration**

Distributed Port Group New Distributed Port Group

Name *

VLAN Type None Standard VLAN

DHCP Service

Cancel **OK**

Figure 23: Distributed Switch

5.2 Distributed Port Group and Kernel Adapters

The next step is to configure the kernel adapters that will allow each host to communicate over the storage network.

A kernel adapter must be created on **each host** in the cluster.

From the distributed switch configuration panel, select the option to create a new **Kernel Adapter**. Following the steps shown in the reference screenshot:

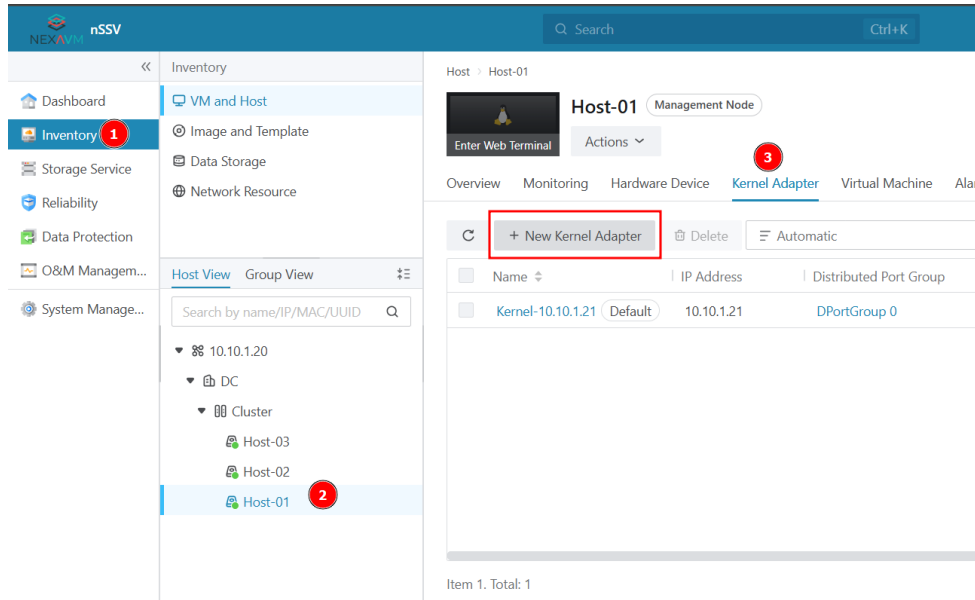


Figure 24: Distributed Switch

- Select the newly created **Distributed Port Group**.
- Assign the **Storage IP** associated with the storage bond on the specific host.
- Enter the correct **Netmask**, as shown in the example screenshot.

New Kernel Adapter | Host-01 ×

Name *

Description
0/256

Network Service

Host Network Configurations

Distributed Port Group * 🔍

IPv4 Address *

Netmask *

Cancel OK

Figure 25: Distributed Switch

Repeat this operation for all cluster hosts to ensure complete connectivity across the storage network.

5.3 Uploading the nSAN Package

With the storage network fully configured, nSAN can now be deployed.

Navigate to:

- **Storage Service** → **Storage Overview**

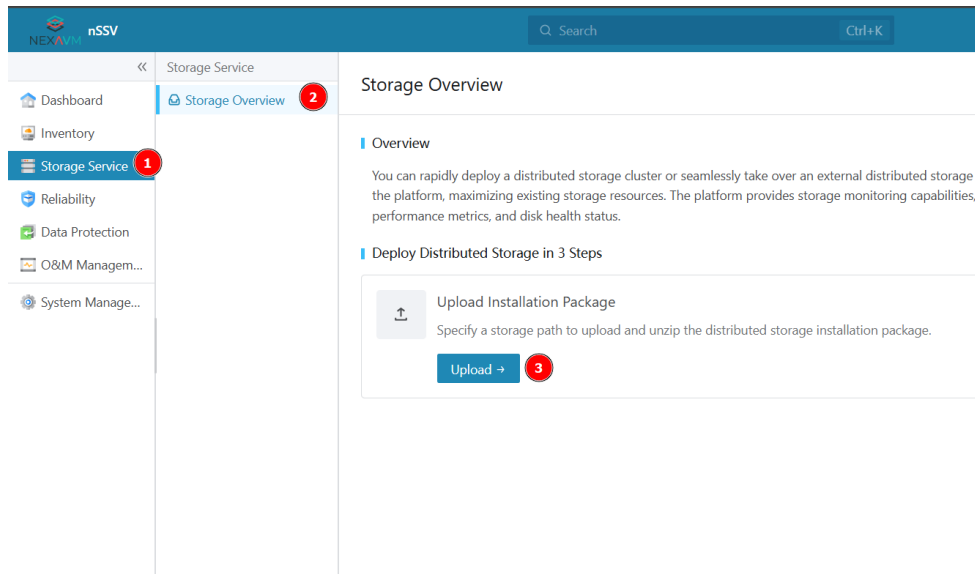


Figure 26: nSAN Package Upload

Click the **Upload** button and select the nSAN installation package.

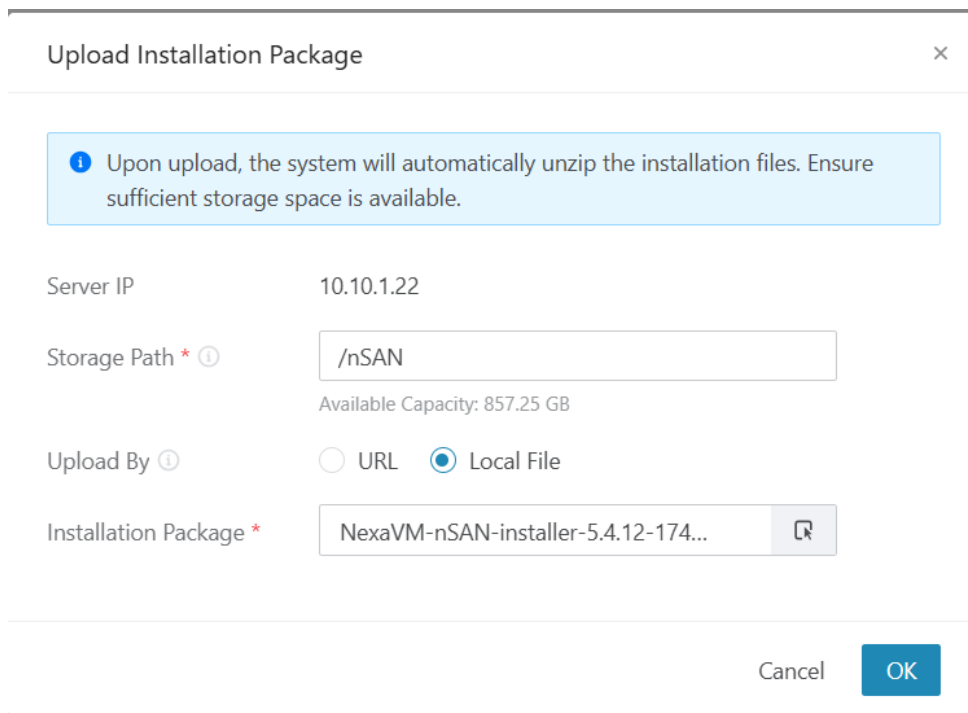





Figure 27: nSAN Package Upload

Follow the installation wizard filling the required fields:

Deploy Management Service×

VIP	10.10.1.20
MN-1 IP	10.10.1.22
SSH Port *	<input type="text" value="22"/>
SSH Password *	<input type="password" value="....."/> 
MN-2 IP	10.10.1.21
SSH Port *	<input type="text" value="22"/>
SSH Password *	<input type="password" value="....."/> 
Database Password	<input type="password" value=""/> 

Cancel OK

Figure 28: nSAN Wizard

During the **Initialize Distributed Storage** phase, as shown in the screenshot below, fill in the fields as follows:

- **Admin Network:** enter the management subnet in use
- **Cluster Network:** enter the storage subnet in use
- **Public Network:** enter the storage subnet in use

Initialize Distributed Storage
×

○
Initialization
●
Storage Configuration

Initialization

Storage MN IP *

Make sure password-free access is configured between all nodes and the hostname of each node is not set to localhost. Otherwise, initialization will fail.

Cluster

Network Configuration

Admin Network *

Used for managing and configuring the storage cluster.

Cluster Network *

Used for monitoring data disks and synchronizing replicas among nodes in the storage cluster.

Public Network *

Used for the storage cluster to provide external services.

Time Sync Server IP *

Ensures time synchronization between all nodes in the storage cluster.

Cancel
Next >

Figure 29: nSAN Wizard

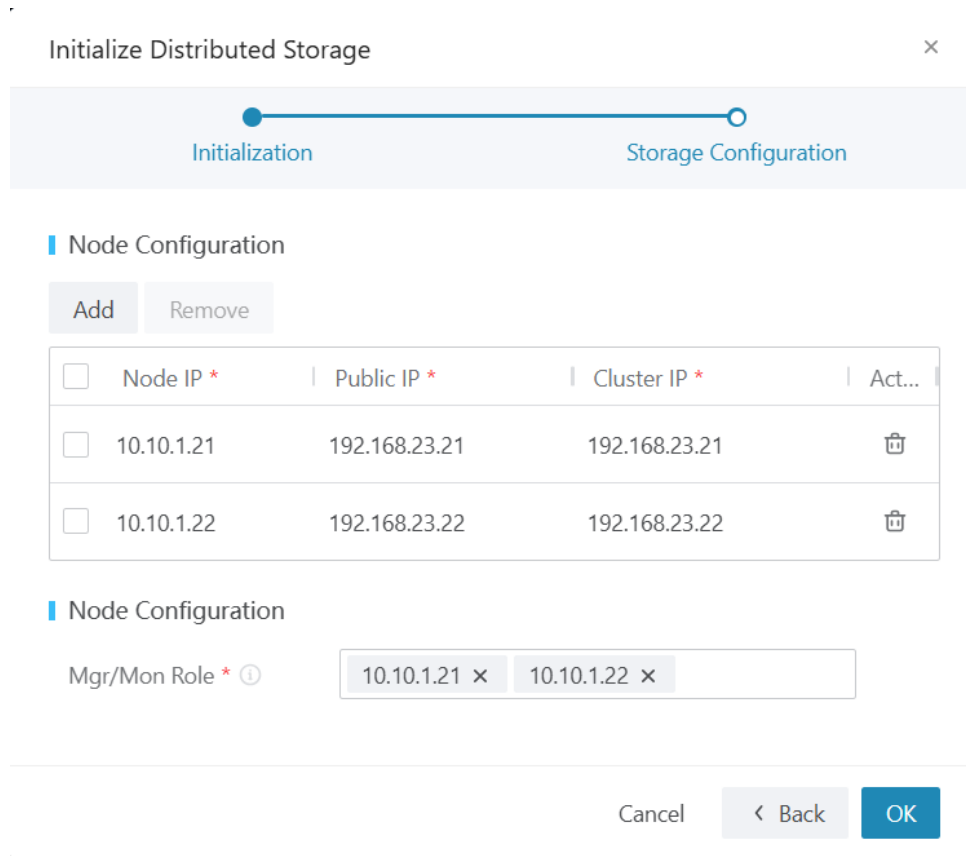


Figure 30: nSAN Wizard

5.4 Adding Storage Nodes to nSAN

After the nSAN package has been uploaded and installation begins, the next step is to register each server that will participate in the distributed storage cluster.

Navigate to:

- **Storage Service** → **Storage Node**

Then click **Add Server** to begin adding the first storage node. Repeat this process for each node that will be used by nSAN.

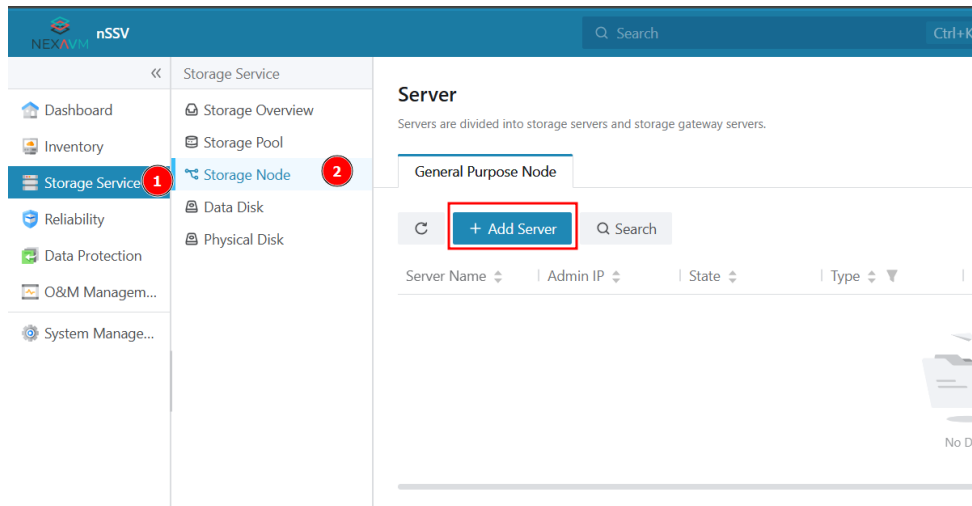


Figure 31: nSAN Wizard

In the “Add Server” dialog, be sure to disable the option **Block Storage Gateway** to prevent configuration of legacy block gateway services that are not required in the hyper-converged setup.

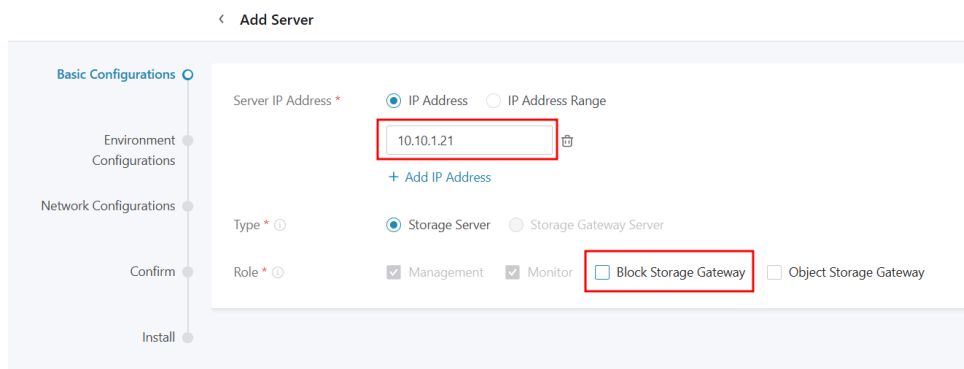


Figure 32: Add Storage Node – Configuration Example

Proceed by entering the hostname or IP address of the node, the root credentials, and any additional parameters requested by the wizard. If the deployment is based on a **2+1 architecture** and the node is intended to act as a **witness**, simply select the node as a monitor during the configuration step. Once the server is added, it will appear in the list of storage nodes and begin synchronizing with the cluster.

The screenshot shows the 'Add Server' configuration page. On the left, a vertical progress bar indicates the current step is 'Basic Configurations'. The main content area contains the following fields and options:

- SSH Username ***: Text input field containing 'root'.
- SSH Password ***: Password input field with masked characters '.....' and a toggle for visibility. Below it, a note states: 'The system uses the password only for password-free login configurations and does not store the password.'
- Port ***: Text input field containing '22'.
- Server Name**: Radio buttons for 'Set Server Name' (selected) and 'Not Set Server Name'. Below is a text input field containing 'host-storage-01'.
- Time Sync Service**: Toggle switch, currently turned on. Below it, a note states: 'Turn on the button to ensure the time consistency of the newly-added servers with other servers in the cluster.'
- Password-Free Login**: Toggle switch, currently turned on. Below it, a note states: 'Turning on the button means that you allow the system to configure password-free logins to the server with the SSH username and password.'

At the bottom right, there are three buttons: 'Cancel', '< Previous', and 'Next >'.

Figure 33: Add Storage Node – Configuration Example

The screenshot shows the 'Add Server' configuration page with the 'Network Configurations' section active. The progress bar on the left shows 'Network Configurations' as the current step. The main content area contains a table for network settings:

Network Configurations			
	Admin IP	Public IP *	Cluster IP *
1	10.10.1.21	192.168.23.21	192.168.23.21

At the bottom left, the progress bar shows 'Confirm' and 'Install' as subsequent steps.

Figure 34: Add Storage Node – Configuration Example

Repeat for all nodes allocated to the nSAN pool to ensure full coverage and redundancy.

5.5 Initializing Physical Disks and Creating Data Disks

After all storage nodes have been added to nSAN, the next step is to prepare the physical disks that will be used to build the virtual storage layer.

From the nSSV web interface, navigate to:

- **Storage Service** → **Physical Disk**

Select the disks that you intend to include in the nSAN storage and click on **Initialize Hard Disk**, as shown in the reference screenshot. This operation prepares the selected devices for use in the distributed storage infrastructure.

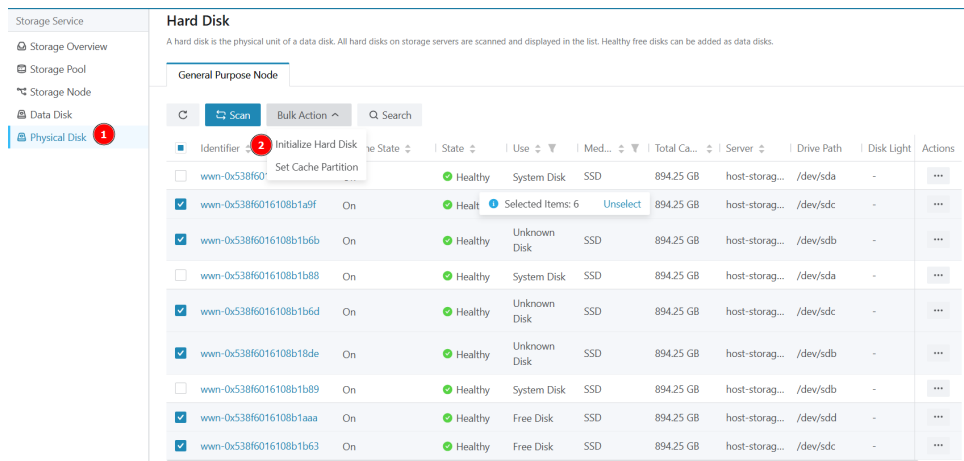


Figure 35: Initializing Physical Disks

Once all required disks have been initialized, proceed to:

- **Storage Service** → **Data Disk**

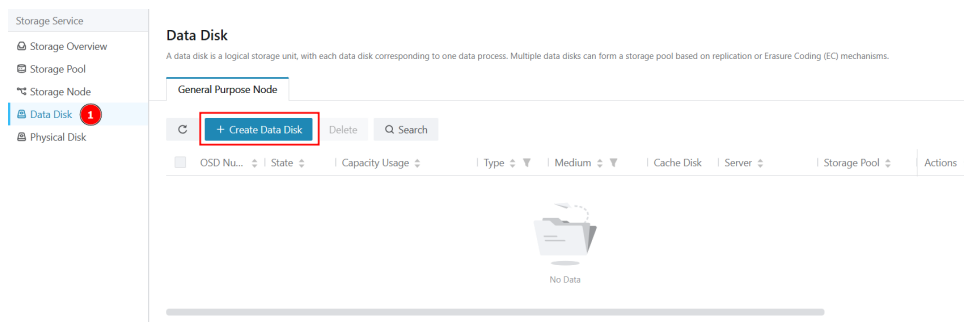


Figure 36: Create Data Disk

Create a new **Data Disk** following the steps illustrated in the example screenshots. In this scenario, the configuration shown is based on a full-flash (all-SSD) environment; therefore, no cache tier was required.

In environments using hybrid disk configurations (SSD + HDD), a cache tier might be necessary, but the procedure in this guide reflects a full-SSD deployment.

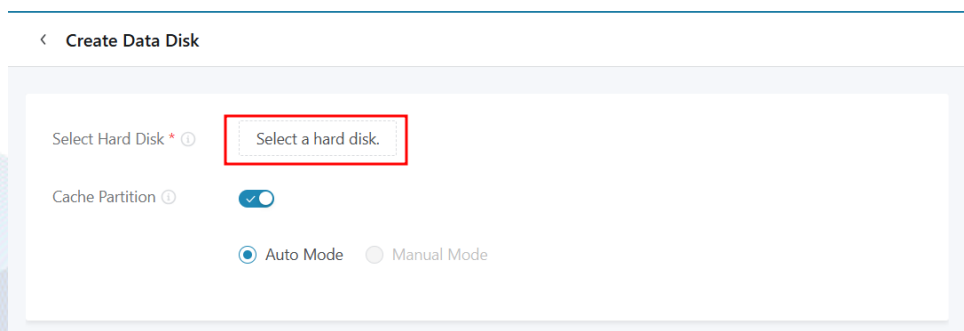


Figure 37: Data Disk Example

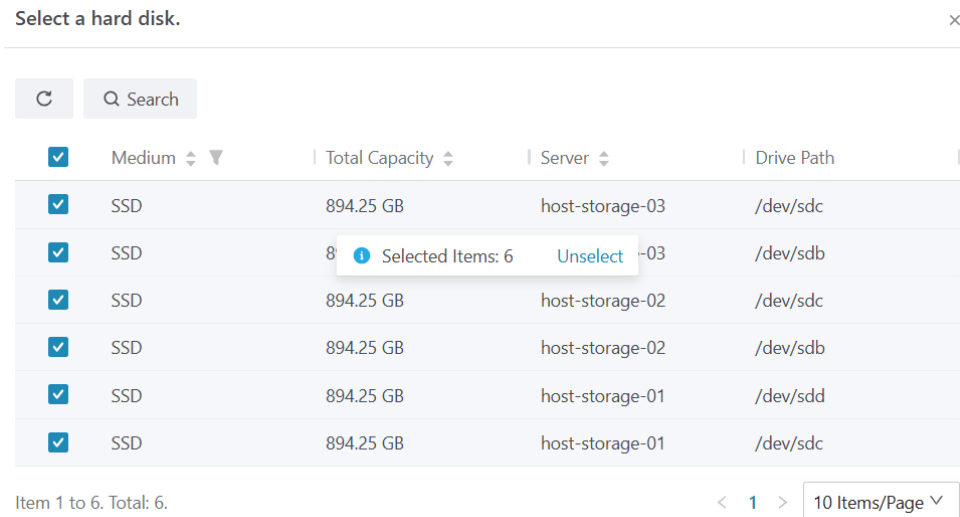


Figure 38: Data Disk Example




Figure 39: Data Disk Example

5.6 Creating the Storage Pool

With the Data Disks prepared, the next step is to create the actual storage pool that will serve as the foundation for the distributed storage layer.

From the nSSV web interface, navigate to:

- **Storage Service** → **Storage Pool**

Click on **Create Storage Pool** to begin the configuration.

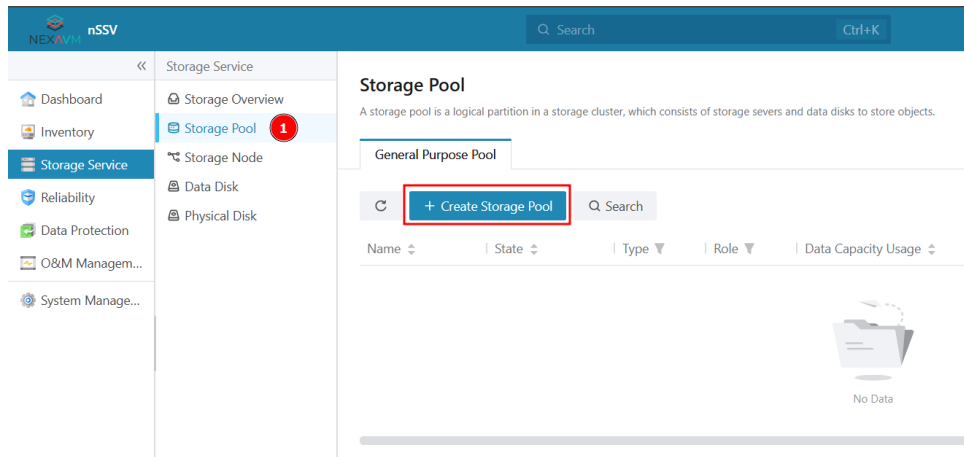


Figure 40: Storage Pool

In the creation wizard:

- Select **Block Storage** as the storage type, since nSAN operates as a block-based distributed storage system.
- Choose the previously created **Data Disk**, ensuring it is assigned to the new storage pool.
- Complete the remaining steps as shown in the example screenshots.

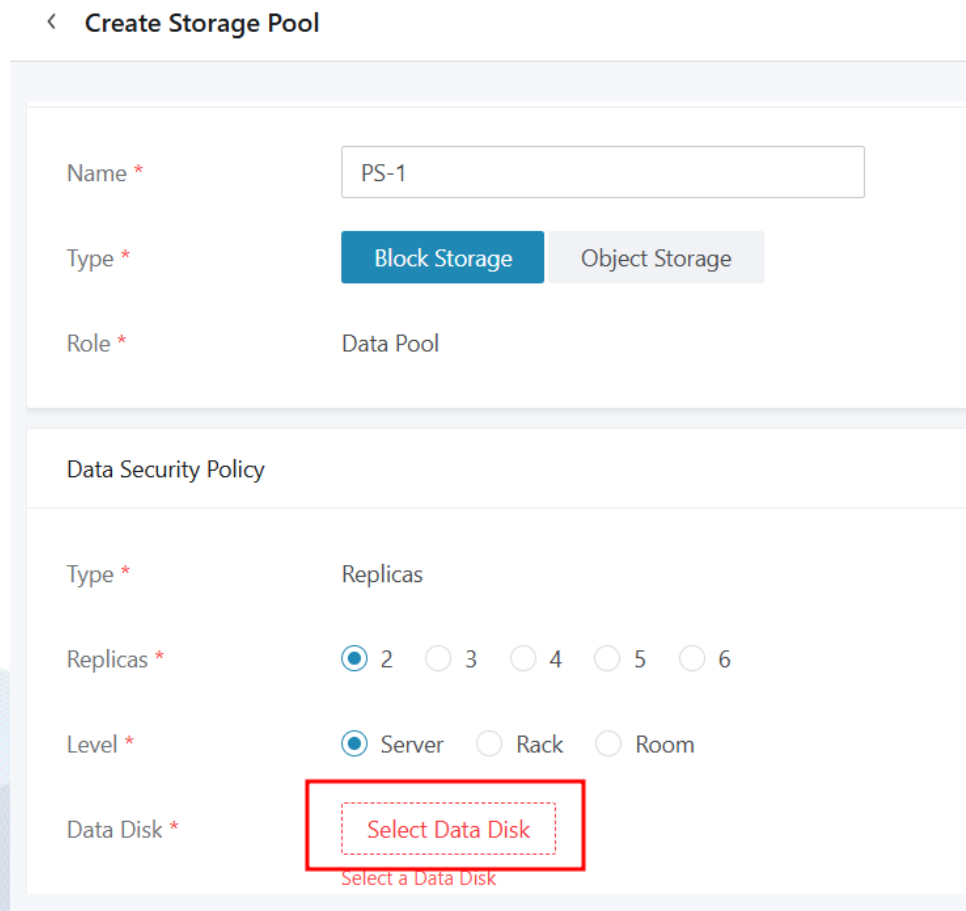


Figure 41: Storage Pool Creation

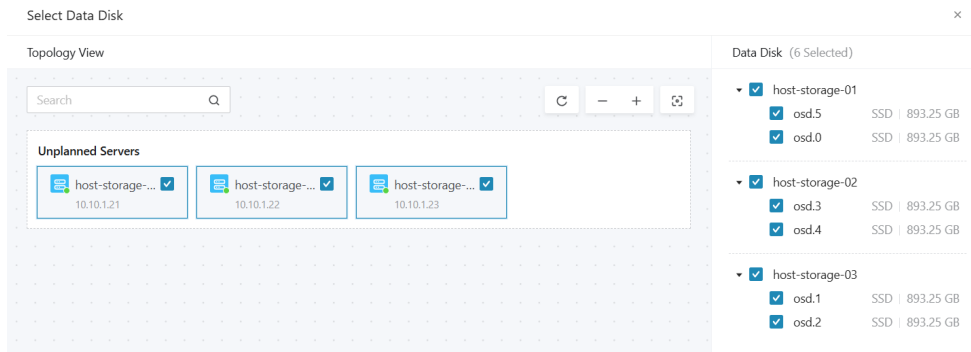


Figure 42: Storage Pool Creation

5.7 Registering the Storage Pool and Creating the Cluster Primary Storage

Once the storage pool has been created, click on its entry in the list to open the detailed view. In this panel, copy the **Pool UUID**, as this value will be required during the creation of the cluster's primary storage.

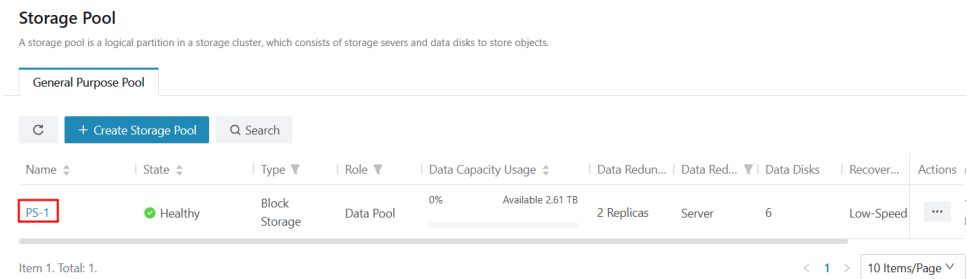


Figure 43: Storage Pool UUID



Basic Info

State : ✔ Healthy

Type : Block Storage

Role : Data Pool

Pool UUID : pool-1553eda824dd423c... Copy

Data Disks : 6

Data Redundan... : Server

QoS Type : StaticQoS (Low Speed)

UUID : 1553eda824dd423cbe739e1...

Figure 44: Storage Pool UUID

After copying the UUID, navigate to:

- **Inventory** → **Data Storage**

Select the previously created **Data Center**, then click on **Actions** and choose **Add Data Storage**.

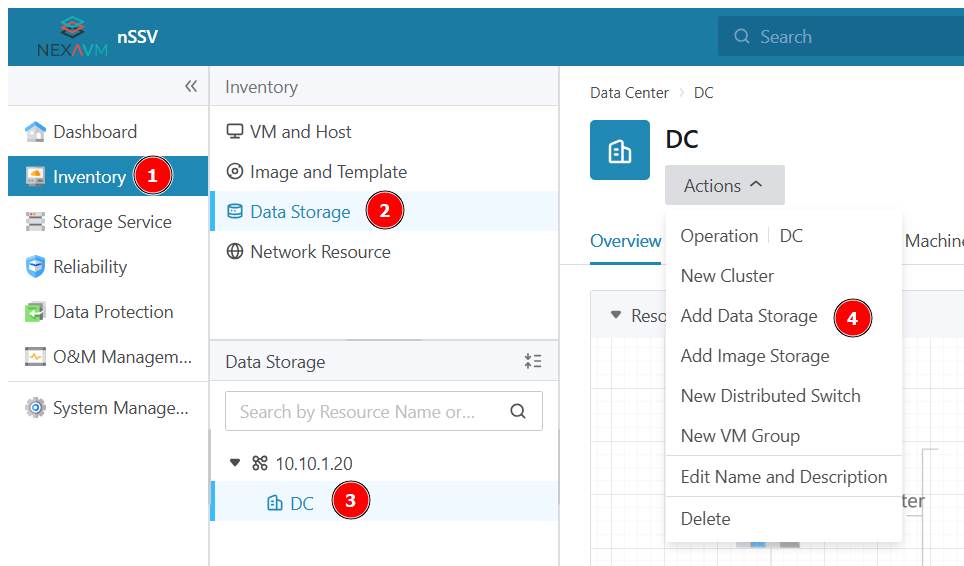


Figure 45: Storage Pool UUID

In the configuration window:

- Select **Distributed Storage nSDS** as the storage type.
- Choose the target **Cluster**.
- Paste the **Pool UUID** into the corresponding field, as shown in the example screenshot.
- Insert the monitoring nodes.

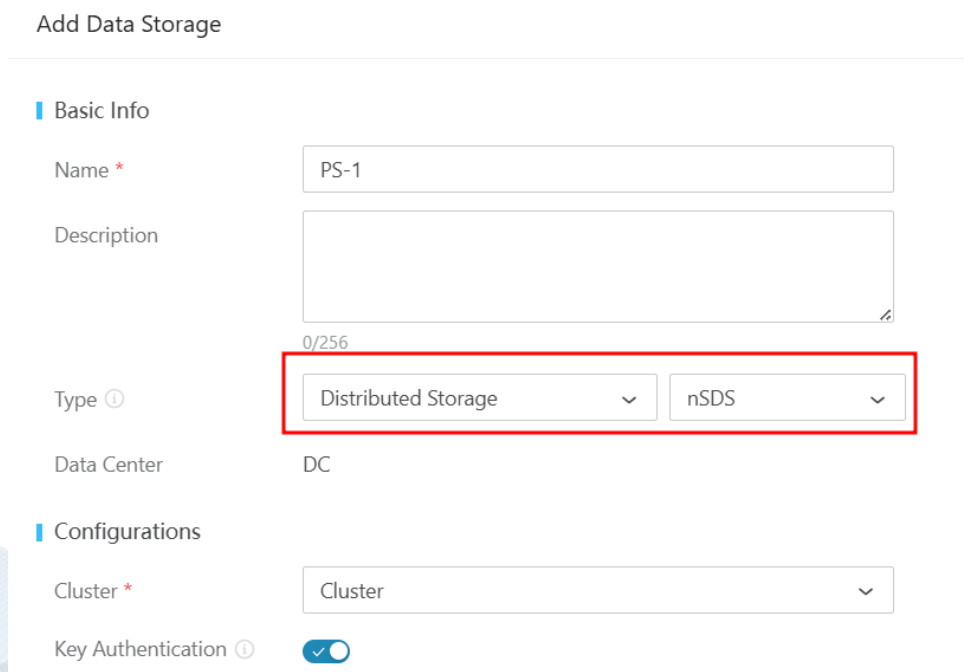


Figure 46: Adding Distributed Storage

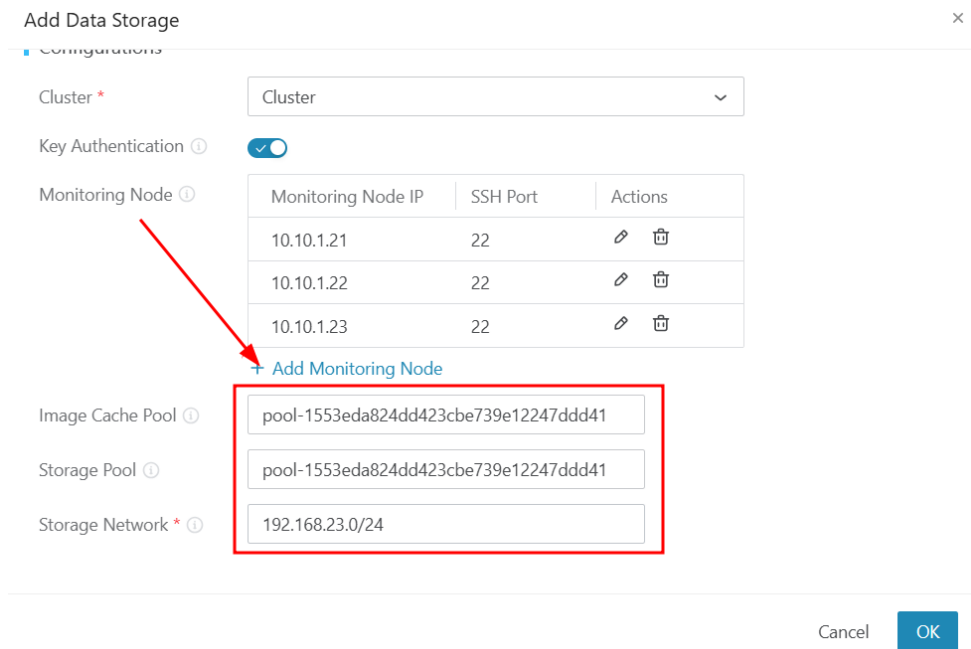


Figure 47: Storage Pool UUID

Once the storage has been created, select it from the list and click on **Attach Cluster** to associate it with the desired cluster.

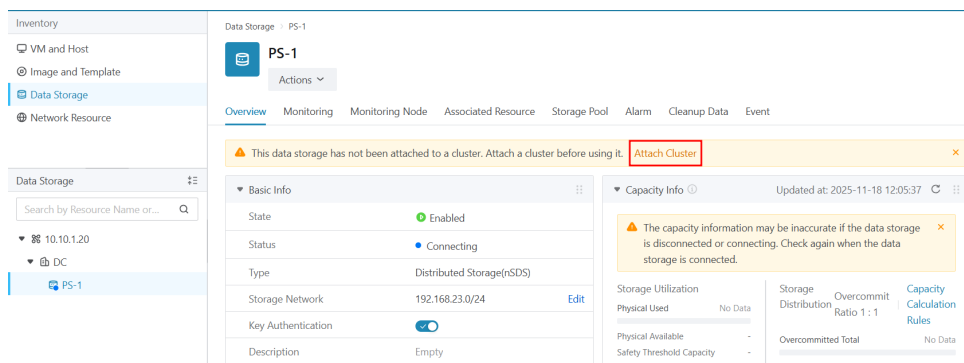


Figure 48: Storage Pool UUID

6 Final Cluster Preparations

After completing the setup of High Availability (HA) and distributed storage (nSAN), the environment requires a few final configurations before virtual machines can be deployed. These steps include creating the Image Storage, uploading OS images, and preparing the virtual network resources used by guest instances.

6.1 Configure Image Storage

To store and manage OS installation media and system templates, an Image Storage must be added to the data center.

1. Navigate to **Image and Template**.

2. Select the appropriate **Data Center** from the left panel.
3. Click **Actions** and choose **Add Image Storage**.
4. Select **Stand Alone Image Storage** as the storage type.
5. Configure the required parameters according to your environment.
6. Click **OK** to create the image storage.

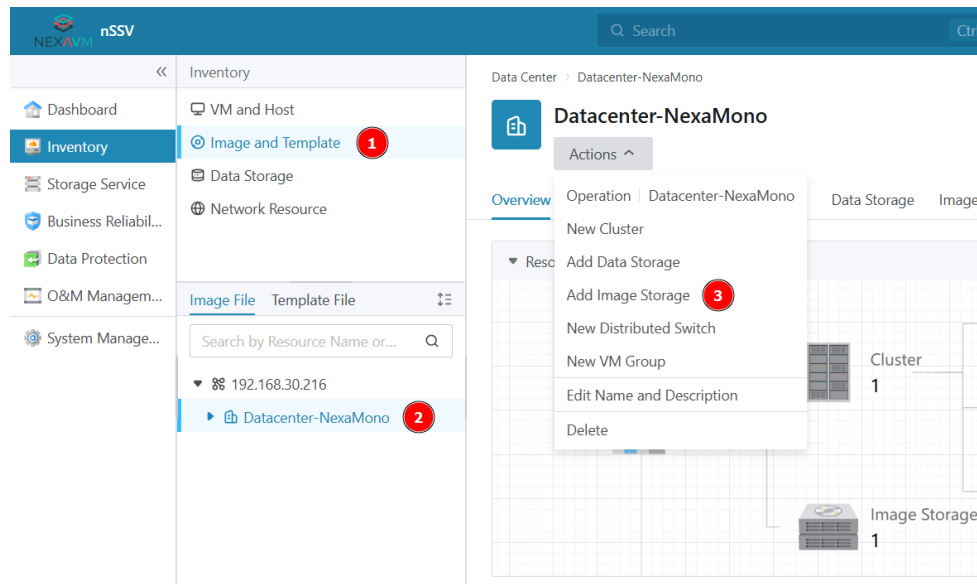


Figure 49: Adding Image Storage

Add OS Images

Once the Image Storage has been created, operating system images can be uploaded.

1. Select the newly created **Image Storage** from the left panel.
2. Click on **Add Image**.
3. Provide the necessary image parameters (name, OS type, format, and source).
4. Click **OK** to begin the upload process.
5. Wait for the upload to complete before proceeding.

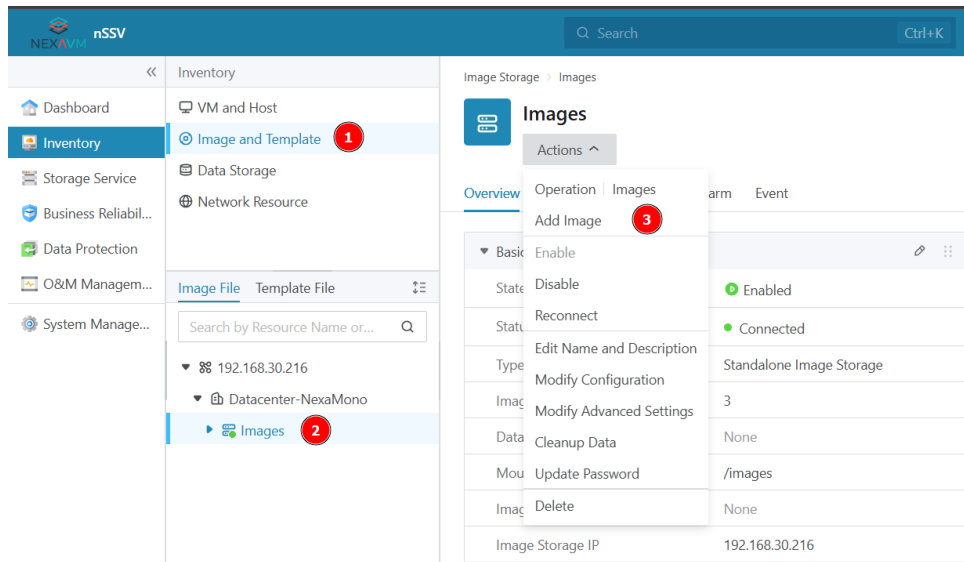


Figure 50: Uploading an OS Image

6.2 Configure Network Resources

Before deploying virtual machines, at least one virtual network must be prepared. If not already done during earlier steps, you can create a Distributed Switch as follows:

1. Navigate to **Network Resources**.
2. Select your **Data Center** from the left panel.
3. Click **Actions** and select **Add New Distributed Switch**.
4. Configure the distributed switch parameters according to your networking design.
5. Click **OK** to finalize the creation.

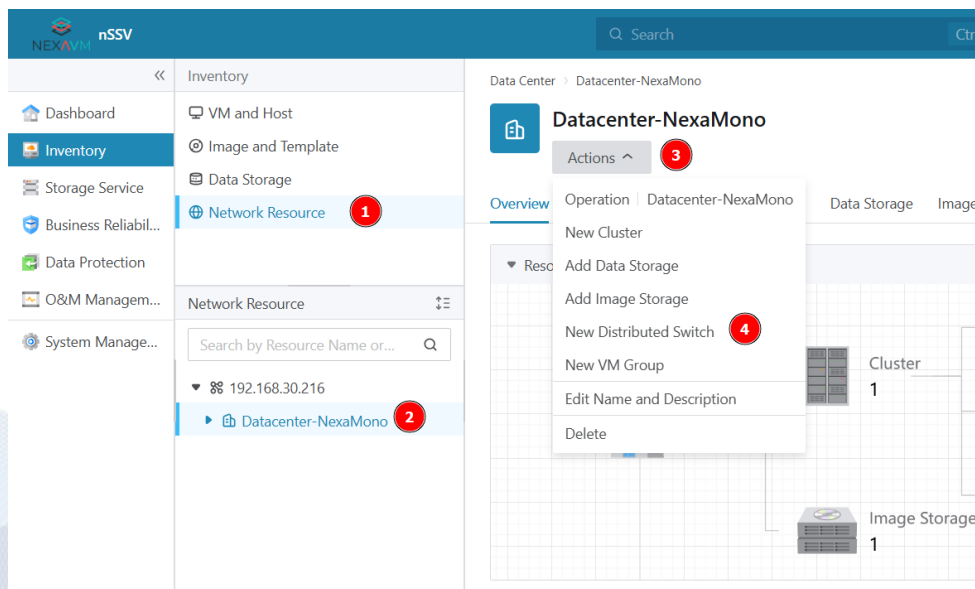


Figure 51: Creating a Distributed Switch

7 Creating Virtual Machines

With the core services of the cluster configured—HA, distributed storage, image storage, and networking—you can now proceed to create virtual machines.

1. Navigate to **VM and Host**.
2. Select the desired **Host** from the left panel.
3. Right-click the host or select the **Actions** menu.
4. Click on **New Virtual Machine**.
5. Configure the VM parameters:
 - **Basic Information:** VM name, description, OS type.
 - **Compute Resources:** CPU and memory allocation.
 - **Storage:** Select the Data Storage and define the disk size.
 - **Networking:** Attach one or more virtual NICs to distributed port groups.
 - **OS Configuration:** Select the installation medium from Image Storage.
6. Review all configuration details and click **OK** to create the virtual machine.




Figure 52: Creating a New Virtual Machine

8 Conclusion of the Installation

At this point, all essential steps for the initial installation of the nHCI platform have been successfully completed. The procedure has guided you through the following phases:

- Installation and synchronization of High Availability and Distributed Storage packages.
- Access to the management interface and execution of the initialization wizard.
- Addition of cluster servers with the appropriate roles depending on the deployment scenario.
- Creation of Data Disks and assembly of the Storage Pool.
- Retrieval and registration of the Storage Pool UUID.
- Configuration of the Primary Storage and association with the nHCI environment.
- License generation and activation.

The system is now operational and ready for use. From this stage onward, administrators can proceed with:

- Deploying and configuring virtual machines.
- Setting up advanced networking features according to infrastructure requirements.
- Integrating monitoring and backup solutions for production environments.
- Applying security policies and best practices to ensure system protection.

Dynamic Expansion: The nHCI platform has been designed to be flexible and scalable. Even after completing the installation described in this guide, it is possible to expand the infrastructure by:

- Adding new compute nodes to increase capacity.
- Integrating additional storage resources into existing storage pools.
- Connecting external storage systems for heterogeneous environments.
- Extending networking features to adapt to future needs.

This completes the base installation procedure. The environment is now ready to be customized and expanded according to the specific requirements of each deployment scenario.