



nSSV Cluster Installation Guide

Clustering with HA and external storage

Contents

1	Overview	2
2	Install nSSV 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	5
2.7	Configure Network Interfaces	6
3	Post-Installation Network Configuration	7
3.1	Login and Initial Setup	7
3.2	Post-Installation Network Configuration	7
3.3	Network Configuration Commands	7
3.4	Removing Incorrect Configurations	8
4	Install nSSV Management Service	9
5	Cluster Initialization	9
5.1	Create Data Center	9
5.2	Create Cluster	10
5.3	Add Host to Cluster	11
6	High Availability Configuration	11
6.1	Accessing the Management Node Operations	11
6.2	Adding the Second Management Node	12
6.3	Accessing the Cluster via the Virtual IP	14
7	External Storage Integration	15
7.1	Adding an iSCSI Storage Target	15
7.2	Creating the External Data Storage	16
8	Final Cluster Preparations	17
8.1	Configure Image Storage	18
8.2	Configure Network Resources	19
9	Creating Virtual Machines	20
10	Conclusion of the Installation	20

1 Overview

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

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)**. 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 nSSV
- Essential hardware and network configuration requirements
- Basic post-installation setup through both CLI and web interface

2 Install nSSV 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/Hyon'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	<p>[leftmargin=*, itemsep=0.2em]</p> <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	–	<p>[leftmargin=*, itemsep=0.2em]</p> <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 nSSV 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 nSSV 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.

Remember to insert your custom root password, it will be used for settings later.

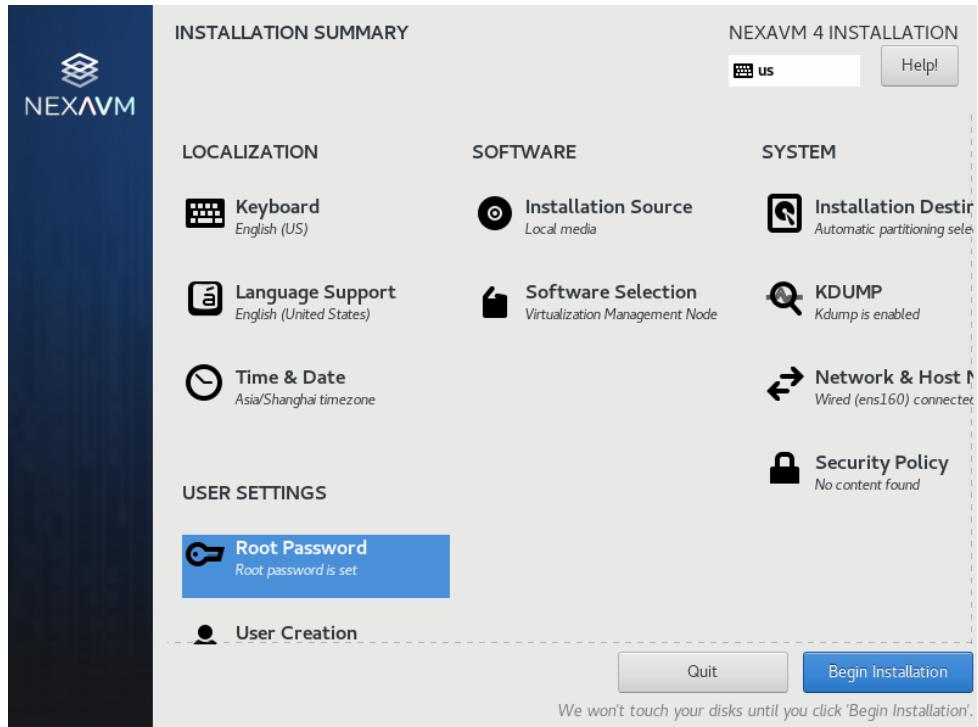


Figure 1: Installation Summary

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 nSSV 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.

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 nSSV (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 Interfaces

On the *Installation Summary* page, click *Network & Host Name* to configure the network interfaces.

- Select the network interface cards (NICs) you want to use
- For each NIC you wish to configure, follow these steps:
 1. Click on the NIC from the list in the left panel

2. Click the **Configure** button in the bottom-right corner
3. In the configuration window, select **IPv4 Settings** from the menu
4. From the **Method** dropdown menu, select **Disabled**
5. Select **IPv6 Settings** from the menu
6. From the **Method** dropdown menu, select **Disabled**
7. Click **Save** to apply the configuration

- Repeat this process for each NIC you want to configure
- After configuring all NICs, click **Done** to return to the main installation screen

This configuration ensures that the network interfaces are properly set up before proceeding with the installation.

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 Post-Installation Network Configuration

3.1 Login and Initial Setup

After the system reboots, log in using the credentials created during the installation process.

3.2 Post-Installation Network Configuration

General Rules

- Replace interface names (e.g., eth0, eth1, eth2, eth3) with the real ones from your system.
- For trunk mode, replace VLAN IDs (100, 200) with your actual VLAN IDs.
- Use the provided IPs, masks, and gateways only as examples.
- Typically, only the management network requires a gateway.

3.3 Network Configuration Commands

The following procedure describes how to configure a bond interface. These steps (`create bond` → `attach NICs` → `optional VLAN` → `bridge and IP configuration`) must be repeated for each bond that the system requires. At a minimum, one bond should be created for **management** and one for **storage**. Additional bonds may be configured as needed for **business traffic**, **backup**, or **migration**.

Create a virtual link aggregation interface in active-backup mode:

```
zs-bond-ab -c [BOND_NAME]
```

Attach a physical NIC to the bond interface:

```
zs-nic-to-bond -a [BOND_NAME] [NIC_NAME]
```

Create a network bridge and configure its IP address:

```
zs-network-setting -b [BOND_NAME] [IP_ADDRESS] [NETMASK] [GATEWAY]
```

(Optional) If required create a VLAN interface, add it to the selected bond, and configure the network:

```
zs-vlan -c [BOND_NAME] [VLAN_ID]
```

```
zs-network-setting -b [BOND_NAME].[VLAN_ID] [NETMASK] [GATEWAY]
```

Verify the current network configuration:

```
zs-show-network
```

3.4 Removing Incorrect Configurations

If a bond or bridge was configured incorrectly:

```
# Stop the created bridge
ip link set [BRIDGE_NAME] down
```

```
# Delete the bridge
brctl delbr [BRIDGE_NAME]
```

```
# Delete the bridge configuration file
rm -f /etc/sysconfig/network-scripts/ifcfg-[BRIDGE_NAME]
```

Delete VLAN configuration:

```
zs-vlan -d [BOND_NAME] [VLAN_ID]
```

Delete bond:

```
zs-bond-ab -d [BOND_NAME]
```

Check that bonds are active and IPs are correctly assigned. Then continue with the remaining nSSV configuration steps.

4 Install nSSV Management Service

After configuring the network, run the following command only on the two management nodes to install the nSSV management service:

```
bash /opt/zstack-installer.bin
```

5 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.

5.1 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

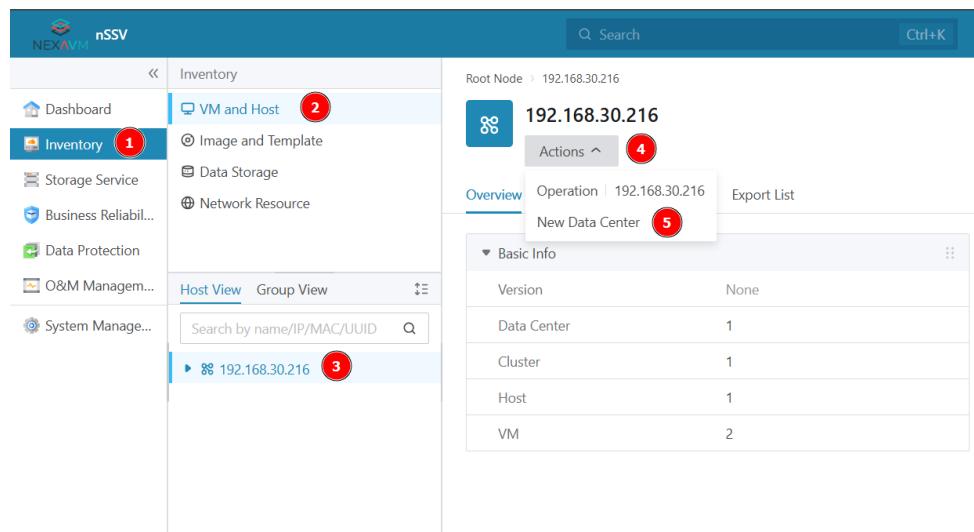


Figure 2: New Data Center Flow

5.2 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

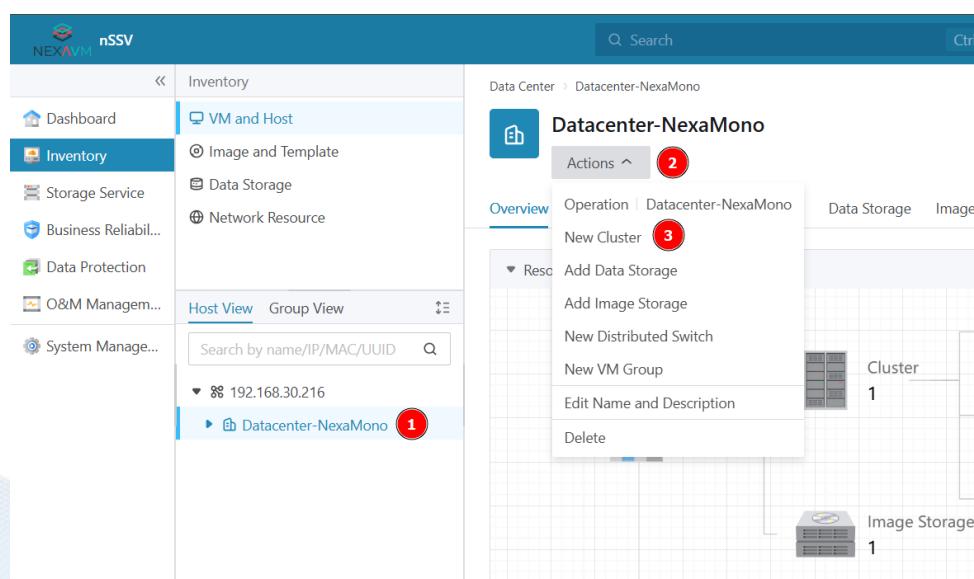


Figure 3: New Data Center Flow

5.3 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 nSSV 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

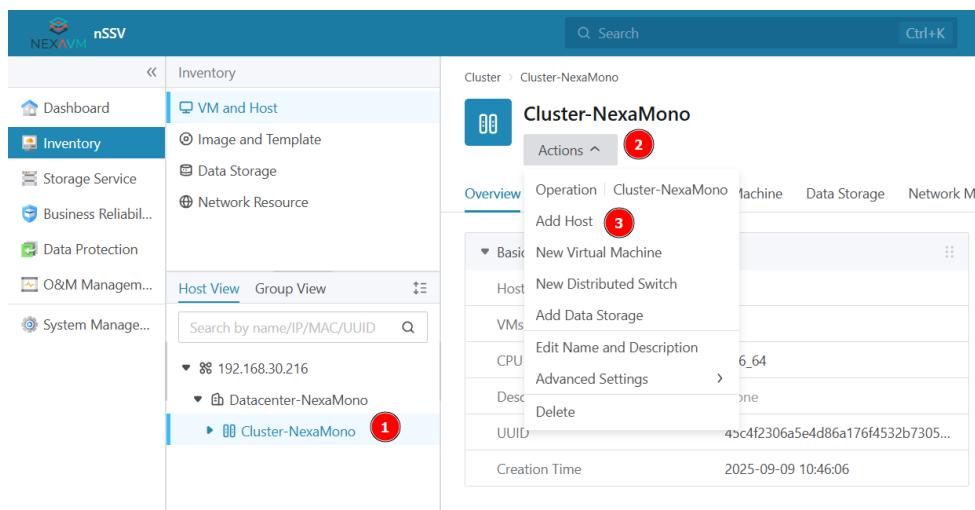


Figure 4: New Data Center Flow

6 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 nSSV environment can continue operating even in the event of a failure of one of the management nodes.

6.1 Accessing the Management Node Operations

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

- **Reliability → MN Monitoring**

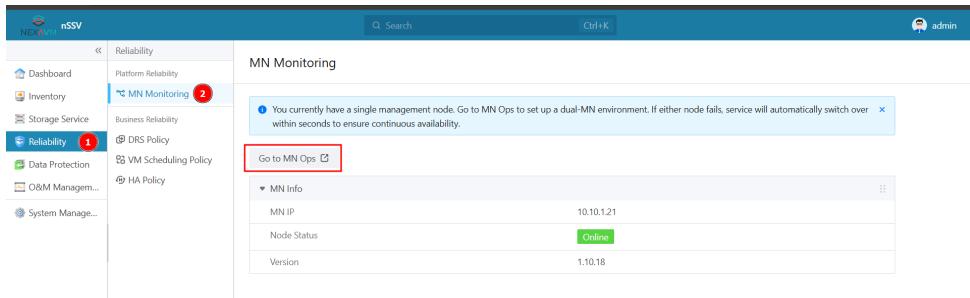


Figure 5: MN Monitoring Navigation

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

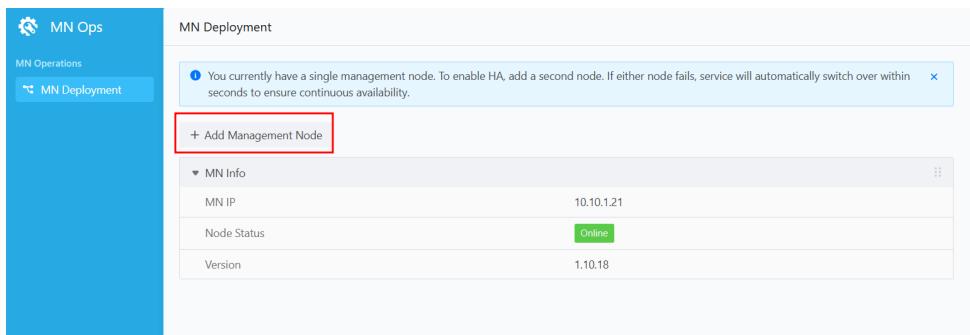


Figure 6: MN Monitoring Navigation

6.2 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.

Add Management Node

Configure MN Review Configuration Add MN

VIP * 10.10.1.20

Active MN 10.10.1.21

SSH Username * root

SSH Password * 

Standby MN IP * 10.10.1.22

SSH Username * root

SSH Password * 

Time Sync Server 10.10.1.21

Cancel Next >

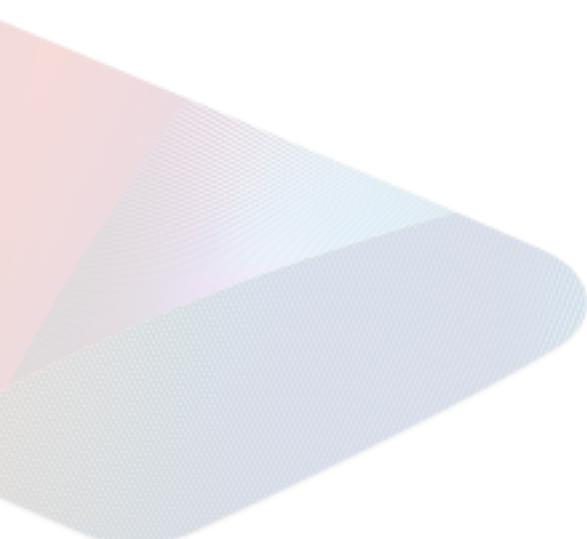


Figure 7: Add Management Node

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

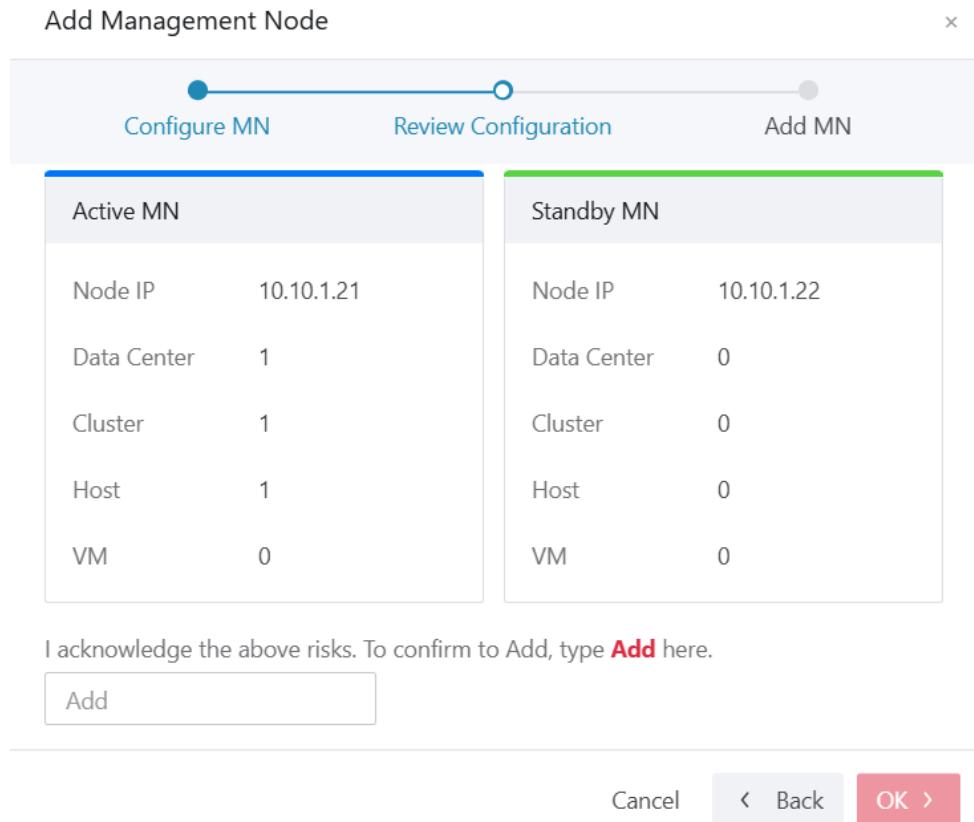


Figure 8: Add Management Node

6.3 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 nSSV management interface using the VIP:

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

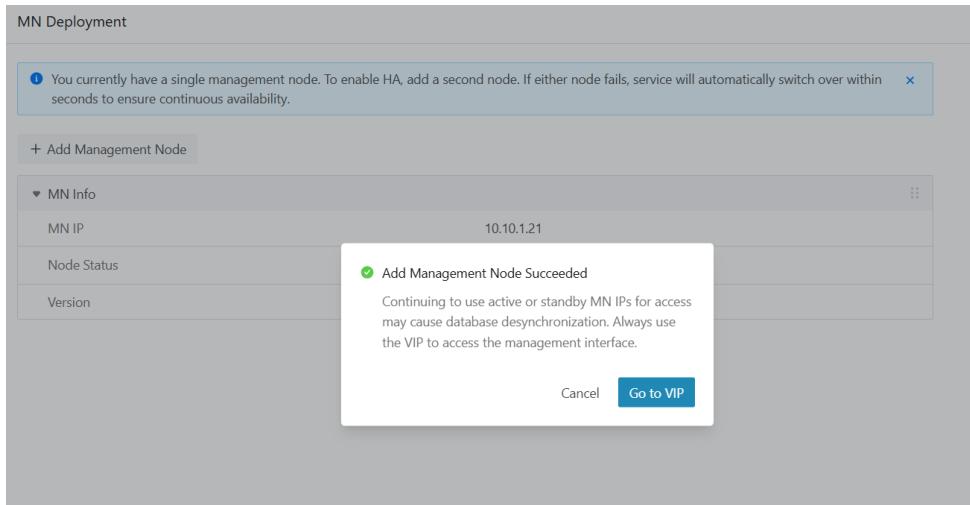


Figure 9: VIP Login

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 nSSV environment.

7 External Storage Integration

In addition to the hyperconverged nSAN storage layer, the platform allows the use of external storage systems. This can be useful in environments where existing SAN infrastructures are present, where specific performance tiers are required, or when migrating workloads from legacy platforms.

In the following example, we illustrate the procedure for integrating an external **iSCSI** storage target. The same workflow applies to other SAN-based technologies, with variations only in the type of target selected.

7.1 Adding an iSCSI Storage Target

To register an external storage device, navigate to:

- **Inventory** → **Data Storage**

Select the previously created **Data Center**, then proceed to:

- **Storage Target**

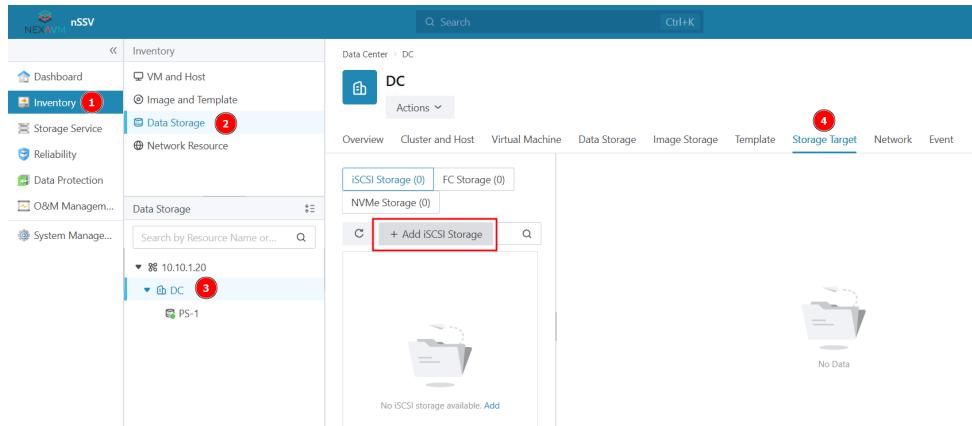


Figure 10: VIP Login

From the list of available storage protocols, select **iSCSI** and click on **Add iSCSI Storage**, as shown in the example screenshot. Enter the required connection parameters provided by your storage system (target IP, port, IQN, authentication details if required).

Add iSCSI Storage

Name *	iSCSI
IP Address *	10.10.1.116
Port *	3260
Cluster *	Cluster
CHAP Username	
CHAP Password	

Cancel OK

Figure 11: Adding an iSCSI Storage Target

7.2 Creating the External Data Storage

Once the iSCSI target has been created, it can be associated with the cluster as a Data Storage resource.

From the same **Data Storage** section:

1. Click **Actions** and select **Add Data Storage**.
2. Choose **SAN** as the storage type.
3. Select the target **Cluster**.

4. Fill in the required network configuration parameters.
5. In the **LUN** section, click on **Add** and select the previously created iSCSI target.

Name	State	Status	Type	Storage Utilization	Available
PS-1	Enabled	Connected	Distributed (nSDS)	0.72%	Available 2.1 TB

Figure 12: Creating a SAN Data Storage and Assigning LUNs

Add Data Storage

Data Center DC

Configurations

Cluster *

Default Provisioning ... * Thick Provision Thin Provision

LUN *

Select LUNs provided by iSCSI, FC, or NVMe storage. Attach an iSCSI storage to the cluster in advance if you need the LUNs it provides.

Storage Network *

Cleanse LUN Cleanse existing data in LUNs

! If data exists in the LUN, you might fail to add LUNs or attach data storage.

Cancel **OK**

Figure 13: Creating a SAN Data Storage and Assigning LUNs

After completing these steps, the external storage system is fully integrated and ready to be used for virtual machine provisioning and other storage operations.

8 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.

8.1 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.

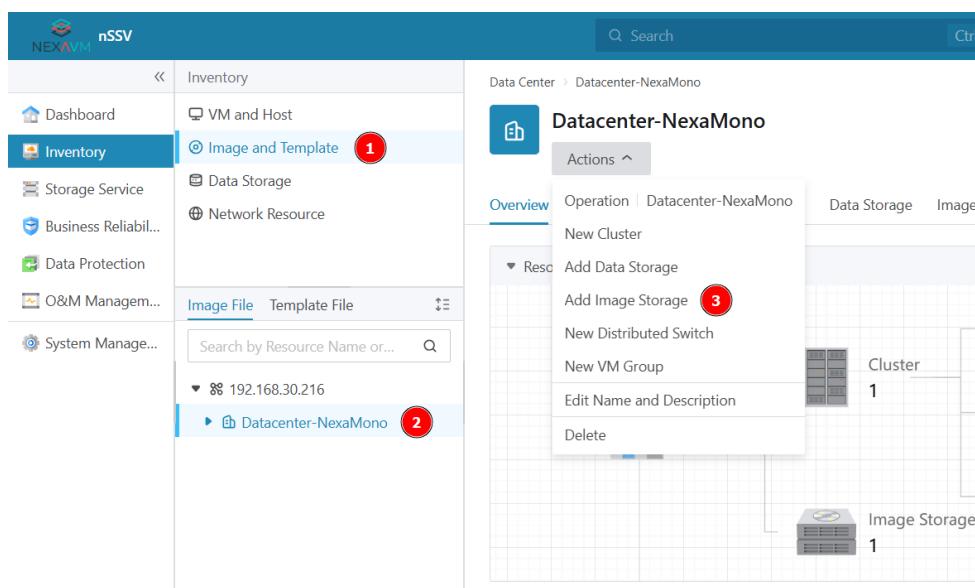


Figure 14: Adding Image Storage

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.

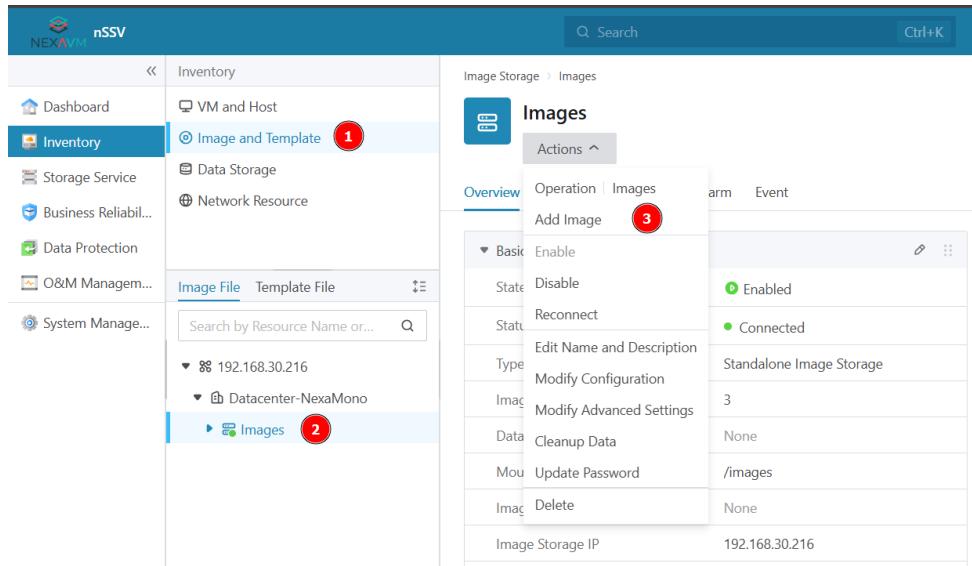


Figure 15: Uploading an OS Image

8.2 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.

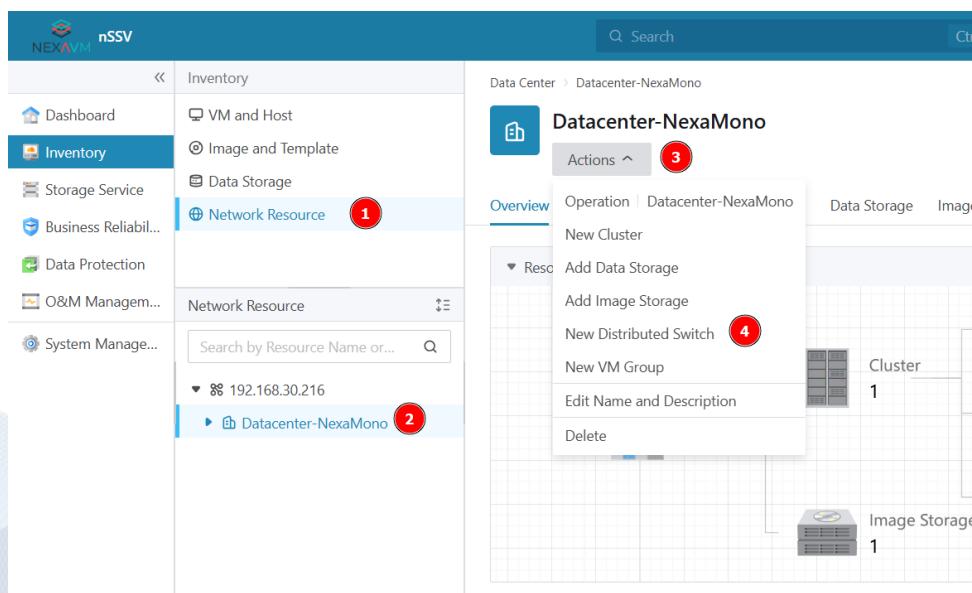


Figure 16: Creating a Distributed Switch

9 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.

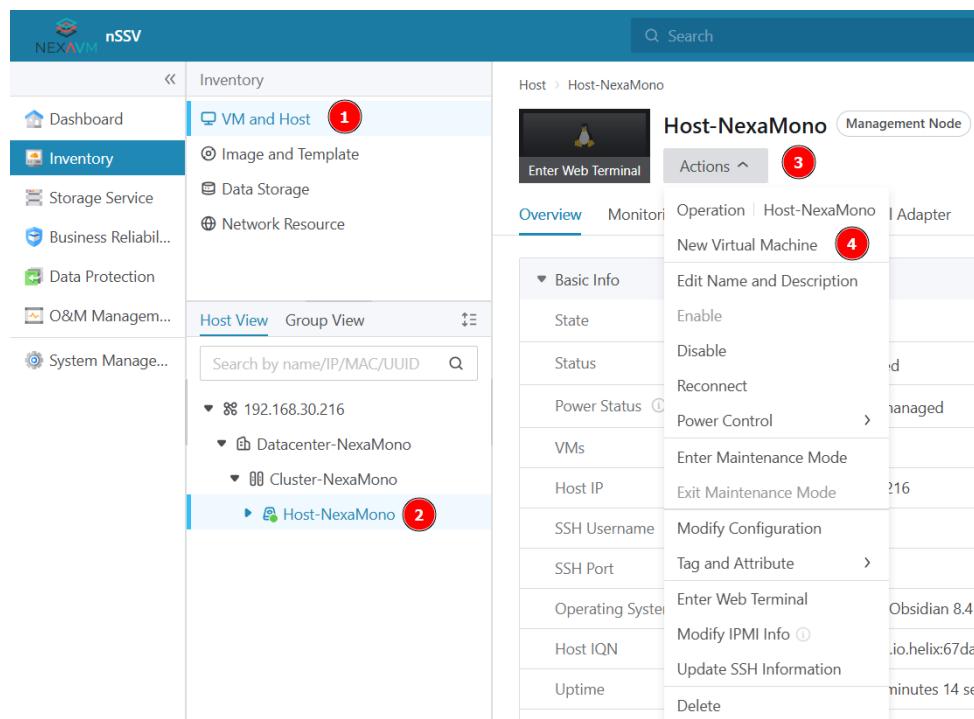


Figure 17: Creating a New Virtual Machine

10 Conclusion of the Installation

At this point, all essential steps for the initial installation of the nSSV 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 nSSV 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 nSSV 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.