NexaSphere

User Guide

Version: V2.1.1

Issue: V2.1.1

Nexavm Technologies AG

NEXAVM

Copyright Statement

Copyright © 2026 Nexavm Technologies AG. All rights reserved.

No part of this document may be excerpted, reproduced, or distributed in any form without prior written permission from Nexavm Technologies AG.

Trademark Notice

NexaSphere and other trademarks of Nexavm Technologies AG are trademarks of Nexavm Technologies AG.

All other trademarks or registered trademarks mentioned in this document are the property of their respective owners.

Notice

The products, services, and features you purchase are governed by the applicable commercial contracts and terms of Nexavm Technologies AG. Some products, services, or features described in this document might not be included in your purchase or available for your use. Unless otherwise agreed in the applicable contract, Nexavm Technologies AG makes no express or implied representations or warranties regarding the contents of this document.

This document might be updated from time to time due to product version upgrades or other reasons. Unless otherwise agreed, this document is provided for guidance only. The statements , information, and recommendations in this document do not constitute any express or implied warranty.

Contents

Copyright Statement..... I
Product Overview..... 1
1.1 What Is NexaSphere?..... 1
1.2 NexaSphere Architecture and Resource Model..... 1
1.2.1 NexaSphere Architecture..... 2
1.2.2 NexaSphere Resource Model..... 2
1.3 NexaSphere User Interface..... 4
Installation and Initialization..... 6
2.1 Overview..... 6
2.2 Planning and Preparation..... 7
2.2.1 Hardware Requirements..... 7
2.2.2 Network Planning..... 8
2.2.3 Prepare Software Packages..... 9
2.2.4 Create Installation Media..... 10
2.2.4.1 Create a USB Drive with Rufus..... 10
2.2.4.2 Create a USB Drive with Fedora Media Writer..... 14
2.2.5 Configure Servers..... 14
2.3 Install the Management Node..... 15
2.3.1 Bonded NIC Deployment..... 15
2.3.2 Single NIC Deployment..... 28
2.3.3 Install with TUI..... 35
2.4 Set Up Management Node HA..... 39
2.4.1 Set Up with the UI..... 40
2.5 Log in to NexaSphere..... 45
2.6 Initialize NexaSphere..... 46
2.6.1 Initialize Manually..... 46
2.6.2 Restore from Backup Data..... 49
Licensing..... 51
3.1 Licensing Modes and License Rules..... 51
3.2 Use a Dedicated License..... 51
3.2.1 Use a USB Key License..... 51
3.2.1.1 Install a USB Key License..... 51
3.2.1.2 Update a USB Key License..... 52
3.2.2 Use a Request Key License..... 52
3.2.2.1 Install a Request Key License..... 52
3.2.2.2 Update a Request Key License..... 53
3.3 View License Status and Licensing Records..... 53
3.4 License Status and Exceptions..... 54
Data Center, Cluster, and Host Management..... 58
4.1 Overview..... 58

4.2 Data Center Management....58
4.2.1 Create a Data Center....58
4.2.2 Edit a Data Center....59
4.2.3 Delete a Data Center....59
4.3 Cluster Management....59
4.3.1 Create a Cluster....59
4.3.2 Modify a Cluster....62
4.3.3 Delete a Cluster....63
4.4 Host Management....63
4.4.1 Add a Host....63
4.4.2 Modify a Host....65
4.4.3 Access a Host....65
4.4.4 Delete a Host....65
4.5 Host Status and Maintenance....66
4.6 Host Hardware Device Management....67
4.6.1 Host NUMA Topology....67
4.6.2 LUN....69
4.6.3 Host NICs....69
4.6.4 GPU Devices....73
4.6.5 USB Devices....74
4.6.6 PCIe Devices....75
Virtual Machine Management....77
5.1 Overview....77
5.2 Prepare Images, Templates, and VM Specifications....77
5.2.1 Image Storage Management....77
5.2.1.1 Add a Standalone Image Storage....78
5.2.1.2 Add a Distributed Image Storage....80
5.2.1.3 Manage Image Storage....81
5.2.1.4 Image Storage Status....82
5.2.2 Image Management....82
5.2.2.1 Add an Image....82
5.2.2.2 Synchronize Images....84
5.2.2.3 Migrate an Image....85
5.2.2.4 Delete an Image....85
5.2.2.5 Export an Image....86
5.2.2.6 Export an OVA Template....86
5.2.3 VM Template Management....87
5.2.3.1 Clone a VM to a Template....87
5.2.3.2 Convert a VM to a Template....88
5.2.3.3 Convert a Template to a VM....89
5.2.3.4 Create a VM from a Template....90
5.2.3.5 Manage VM Templates....93
5.2.4 VM Specification Management....94
5.2.4.1 Create a Windows VM Specification....95

5.2.4.2 Create a Linux VM Specification..... 97
5.2.4.3 Manage VM Specifications..... 98
5.3 Deploy VM and Manage the VM Lifecycle..... 99
5.3.1 Create a Virtual Machine..... 99
5.3.2 Import a Virtual Machine..... 110
5.3.3 Register a Virtual Machine..... 112
5.3.4 Clone a Virtual Machine..... 112
5.3.5 Delete a Virtual Machine..... 116
5.4 Access, Group, and Power Management..... 117
5.4.1 VM Group Management..... 117
5.4.2 Access a VM..... 118
5.4.2.1 Manage VM Access and Boot Options..... 118
5.4.2.1.1 Configure Boot Options..... 118
5.4.2.1.2 Configure Remote Access..... 120
5.4.2.1.3 Configure Login Authentication..... 121
5.4.2.2 Access a VM by Using the Console..... 122
5.4.2.3 Access a VM by Using SSH..... 123
5.4.2.4 Console Proxy Management..... 123
5.4.3 VM Power Management..... 124
5.5 Install and Use VMTools..... 127
5.5.1 Install VMTools on a Linux VM..... 129
5.5.2 Install VMTools on a Windows VM..... 130
5.5.3 VMTools Supported OS List..... 131
5.6 Configure a VM..... 135
5.6.1 Modify VM Encryption Settings..... 135
5.6.2 VM Time Synchronization..... 136
5.6.3 Import User Data..... 137
5.6.3.1 Import User Data to a Linux VM..... 137
5.6.3.2 Import User Data to a Windows VM..... 138
5.6.4 Configure VM Resource Contention..... 139
5.6.5 Configure VM CPU Features..... 142
5.6.6 Configure VM QoS..... 143
5.7 Migrate VMs..... 144
5.7.1 Change Host..... 144
5.7.1.1 Hot Migrate a VM to Another Host..... 144
5.7.1.2 Cold Migrate a VM to Another Host..... 145
5.7.2 Change Data Storage..... 146
5.7.2.1 Hot Migrate a VM from SAN Storage to SAN Storage..... 147
5.7.2.2 Cold Migrate a VM from SAN Storage to SAN Storage..... 148
5.7.2.3 Hot Migrate a VM Between SAN Storage and nSDS Distributed Storage..... 150
5.7.2.4 Migrate Data Disks..... 151
5.7.3 Change Host and Data Storage..... 152
5.7.3.1 Hot Migrate a VM Across Data Storage of the Same Type..... 152
5.7.3.2 Cold Migrate a VM Across Data Storage of the Same Type..... 154

5.7.3.3 Cold Migrate a VM Across Different Data Storage Types..... 155
5.7.3.4 Hot Migrate a VM Across Storage Pools in the Same nSDS Distributed Storage.....156
5.7.3.5 Cold Migrate a VM Across Storage Pools in the Same nSDS Distributed Storage..... 157
5.8 Manage VMs with Snapshots..... 158
5.8.1 Create a Snapshot.....159
5.8.2 Revert to a Snapshot.....160
5.8.3 Create a VM from a Snapshot.....161
5.8.4 View Snapshots..... 163
5.8.5 Delete a Snapshot..... 163
5.8.6 Snapshot Policy Management..... 164
5.8.6.1 Create a Snapshot Policy..... 164
5.8.6.2 Enable or Disable a Snapshot Policy..... 165
5.8.6.3 Modify a Snapshot Policy..... 165
5.8.6.4 Delete a Snapshot Policy..... 165
5.8.7 Snapshot Best Practices.....166
5.9 Maintenance and Troubleshooting.....166
5.9.1 Configure a VM Failure Response Policy.....167
Storage.....168
6.1 Overview.....168
6.2 Data Storage Types and Resource Relationships.....168
6.2.1 Data Storage Types.....168
6.2.2 Storage Devices and SAN Storage Targets.....169
6.2.3 Cluster, Data Storage, and Image Storage Relationships.....169
6.3 SAN Storage Target Management.....171
6.3.1 Add an iSCSI Storage Target.....171
6.3.2 Synchronize an FC Storage Target.....172
6.3.3 Add an NVMe Storage Target.....172
6.4 Add Data Storage..... 173
6.4.1 Add Local Storage..... 173
6.4.2 Add NFS Storage.....174
6.4.3 Add SAN Storage..... 176
6.4.4 Add nSDS Distributed Storage.....178
6.4.5 Add ZHPS Distributed Storage.....179
6.4.6 Add ZBS Distributed Storage..... 180
6.5 State, Status, and Maintenance.....181
6.6 Configuration..... 182
6.6.1 Modify Local Storage Configuration.....182
6.6.2 Modify NFS Storage Configuration.....183
6.6.3 Modify SAN Storage Configuration.....184
6.6.4 Distributed Storage Configuration Management.....185
6.7 Capacity and Space Management.....187
6.7.1 Expand SAN Storage.....187

6.7.2 Clean Up Data Storage..... 188
6.8 Delete Data Storage..... 188
6.9 Distributed Storage Service Deployment and Takeover..... 188
6.9.1 Deploy Distributed Storage in 3 Steps..... 189
6.9.1.1 Step 1: Upload the Installation Package..... 189
6.9.1.2 Step 2: Deploy the Management Service..... 191
6.9.1.3 Step 3: Initialize Distributed Storage..... 191
6.9.2 Take Over Existing Distributed Storage..... 193
6.9.2.1 Take Over Distributed Storage..... 193
6.9.2.2 Cancel Distributed Storage Takeover..... 194
Networking..... 195
7.1 Overview..... 195
7.2 Network Model and Resources..... 195
7.3 Configure Distributed Switches..... 196
7.3.1 Create a Distributed Switch..... 196
7.3.2 Attach or Detach Clusters..... 201
7.3.3 Delete a Distributed Switch..... 202
7.4 Manage Uplinks..... 202
7.4.1 Modify Uplink Configuration..... 202
7.4.2 Manage Joined Hosts..... 203
7.4.3 Configure Uplinks for Hosts..... 204
7.5 Configure Distributed Port Groups..... 204
7.5.1 Create a Distributed Port Group..... 204
7.5.2 Modify Distributed Port Group Configuration..... 207
7.5.3 Delete a Distributed Port Group..... 207
7.6 Configure Kernel Adapters..... 208
7.6.1 Create a Kernel Adapter..... 208
7.6.2 Modify Kernel Adapter Configuration..... 208
7.6.3 Delete a Kernel Adapter..... 209
7.7 View Network Topology..... 209
7.7.1 Supported Network Topology Operations..... 210
7.8 Manage Security Groups..... 210
7.8.1 Create Security Groups and Rules..... 212
7.8.2 Operate Security Groups and Rules..... 215
Resource Management and Availability..... 217
8.1 Overview..... 217
8.2 Resource Allocation and Overcommit..... 217
8.2.1 Configure Cluster Overcommit..... 218
8.2.2 Configure Storage Overcommit..... 219
8.2.3 Resource Allocation and Overcommit Formulas..... 220
8.3 Cluster DRS..... 221
8.3.1 Configure DRS..... 221
8.3.2 Perform DRS Operations..... 223

8.4 VM Scheduling Policy..... 223
8.4.1 Create Mutually Exclusive/Affinitive VM Scheduling Policies.....225
8.4.2 Create VMs Exclusive from/Affinitive to Hosts Scheduling Policies.....226
8.4.3 Manage VM Scheduling Policies and Related Resources.....228
8.5 VM HA Policy..... 229
8.5.1 Configure HA Policy.....230
8.5.2 View HA Tasks..... 233
8.5.3 Disable HA Policy..... 233
8.5.4 HA Policy Practice..... 234
Monitoring and Operations.....235
9.1 Overview.....235
9.2 Resource Overview and Capacity Monitoring.....235
9.2.1 View Dashboard Resource and Alarm Statistics..... 235
9.2.2 View Data Center Resource Topology..... 237
9.2.3 View Resource Capacity Information.....237
9.3 Resource Performance Monitoring..... 238
9.3.1 View Monitoring Charts.....239
9.3.2 Customize Monitoring Charts.....239
9.3.3 Metrics Reference..... 240
9.4 Platform and Hardware Health Monitoring.....243
9.4.1 View Management Node Monitoring.....243
9.4.2 View Host Hardware Monitoring..... 244
9.5 Alarms and Notifications..... 244
9.5.1 Prepare the Alarm Service..... 245
9.5.2 Endpoints..... 246
9.5.2.1 Create an Email Endpoint.....246
9.5.2.2 Create a DingTalk Endpoint..... 246
9.5.2.3 Create a Lark Endpoint.....247
9.5.2.4 Create a WeCom Endpoint.....248
9.5.2.5 Create an SMS Endpoint.....249
9.5.2.6 Create an HTTP Application Endpoint.....250
9.5.2.7 Create a Microsoft Teams Endpoint.....250
9.5.2.8 Create an SNMP Trap Receiver Endpoint.....251
9.5.2.9 Manage Endpoints..... 252
9.5.3 Message Templates.....252
9.5.3.1 Create a Message Template..... 252
9.5.3.2 Manage Message Templates.....254
9.5.4 Alarms..... 255
9.5.4.1 Alarm Rules..... 255
9.5.4.2 Create a Resource Alarm..... 256
9.5.4.3 Create an Event Alarm..... 257
9.5.4.4 Manage Alarms..... 257
9.5.5 Alarm Messages..... 258
9.5.5.1 View Alarm Messages..... 258

9.5.5.2 Acknowledge Alarm Messages....258
9.5.5.3 Set a Silence Period....259
9.5.5.4 Restore Alarm Messages....259
9.6 Tasks, Events, and Audit....259
9.6.1 View Tasks....259
9.6.2 View Events....260
9.7 Log Collection and Management....260
9.7.1 Collect Logs....261
9.7.2 Manage Collected Logs....261
9.8 SNMP Monitoring and Notification Integration....262
9.8.1 Configure SNMP Management....262
9.8.2 Download MIB Files....264
9.8.3 Manage SNMP Trap Receivers....264
9.9 Resource Tagging and Search....264
9.9.1 Manage Tags....265
9.9.2 Search Resources by Tag....266
9.9.3 Create a Custom Attribute....266
9.9.4 Add Custom Attribute Values to Resources....267
9.9.5 Manage a Custom Attribute....268
9.9.6 Search Resources by Custom Attribute....268
Security....269
10.1 Overview....269
10.2 Identity and Access Management....269
10.2.1 Prerequisites....269
10.2.2 User Management....269
10.2.2.1 Create a User....270
10.2.2.2 Disable or Enable a User....271
10.2.2.3 Modify User Configuration....271
10.2.2.4 Change a User Password....272
10.2.2.5 Convert a User to an Admin User....272
10.2.2.6 Delete a User....272
10.2.3 User Group Management....273
10.2.3.1 Create a User Group....273
10.2.3.2 Modify User Group Configuration....274
10.2.3.3 Delete a User Group....274
10.2.4 Role and Permission Management....275
10.2.4.1 Built-in Roles....275
10.2.4.2 Create a Custom Role....275
10.2.4.3 Clone a Role....276
10.2.4.4 Modify Role Permissions....276
10.2.4.5 Delete a Role....277
10.2.5 Single Sign-On....277
10.2.5.1 Add an OIDC SSO Server....278
10.2.5.2 Add an AD SSO Server....279

10.2.5.3 Add an LDAP SSO Server..... 281
10.2.5.4 Manage SSO Servers..... 283
10.3 Platform Access Security..... 283
10.3.1 IP Allowlist and Blocklist..... 283
10.3.1.1 Add an IP Allowlist or Blocklist Rule..... 284
10.3.1.2 Manage IP Allowlist and Blocklist Rules..... 284
10.3.2 Security Settings..... 284
10.3.2.1 Modify Security Settings..... 285
10.3.2.2 Security Setting Options..... 285
10.4 SSL Certificate..... 286
10.4.1 Import a Third-Party Certificate..... 287
10.4.2 Import a System Self-Signed Certificate..... 288
10.4.3 Update a Certificate..... 290
10.4.4 Switch to HTTP Login for the UI..... 290
10.5 AccessKey Management..... 291
10.5.1 Generate an AccessKey..... 291
10.5.2 Manage AccessKeys..... 291
10.6 Key Provider Management..... 291
10.6.1 Configure and Manage a Standard Key Provider..... 292
10.6.1.1 Add a Standard Key Provider..... 292
10.6.1.2 Establish Trust..... 293
10.6.1.2.1 Use KMS Certificate and Private Key to Make KMS Trust NexaSphere.. 293
10.6.1.2.2 Use the New Certificate Signing Request (CSR) to Make KMS Trust NexaSphere..... 294
10.6.1.2.3 Make NexaSphere Trust the KMS..... 294
10.6.1.3 Delete a Standard Key Provider..... 295
10.6.1.4 Use the Thales Key Provider..... 295
10.6.1.5 Use the HashiCorp Key Provider..... 307
10.6.2 Configure and Manage a Native Key Provider..... 311
10.6.2.1 Add a Native Key Provider..... 311
10.6.2.2 Back Up the Native Key Provider..... 311
10.6.2.3 Restore the Native Key Provider..... 313
10.6.2.4 Delete the Native Key Provider..... 313
10.6.3 Rekey..... 314
10.7 VM Encryption..... 314
10.7.1 VM Encryption and Disk Encryption..... 315
10.7.2 Encryption Keys and How Encryption Works..... 315
10.7.3 Encryption Requirements and Limitations..... 316
10.8 VM Trusted Security..... 319
10.8.1 Secure VMs with vTPM..... 319
Data Protection and Recovery..... 321
11.1 Overview..... 321
11.2 Prerequisites..... 321
11.3 Add Backup Storage..... 322

11.3.1 Add Local Backup Storage....322
11.3.1.1 Reuse Image Storage....323
11.3.1.2 Reuse Host....323
11.3.1.3 Dedicated Backup Storage....325
11.3.2 Add Remote Backup Storage....326
11.3.2.1 Dedicated Backup Storage....327
11.4 Virtual Machine Backup....328
11.4.1 Backup Policy....329
11.4.2 Create a Backup Plan....330
11.4.3 View Backup Plans and Data....333
11.4.3.1 View Backup Plan Executions....333
11.4.3.2 View VM Backup Data....334
11.4.4 Manual Backup....335
11.4.4.1 Run a Backup Now....335
11.4.4.2 Create an On-Demand Backup....335
11.5 Platform Database Backup....336
11.5.1 Backup Policy....336
11.5.2 Create a Backup Plan....337
11.5.3 View Backup Plans and Data....339
11.5.3.1 View Backup Plan Executions....340
11.5.3.2 View Platform Database Backup Data....340
11.6 Data Recovery....341
11.6.1 Restore a VM from Backup....341
11.6.2 Create a VM from Backup....342
11.6.3 Restore the Platform Database from Backup....343
11.7 Backup Data Management....344
11.8 Backup Plan Management....346
11.9 Backup Storage Management....348
Bare Metal Management....351
12.1 Overview....351
12.2 Preparation....352
12.3 Quick Start Guide....355
12.4 Bare Metal Templates....355
12.4.1 Template Syntax Rules....355
12.4.2 System Preconfigured Templates....356
12.4.3 Add a Custom Template....357
12.5 Bare Metal Clusters....358
12.5.1 Create a Bare Metal Cluster....358
12.5.2 Manage Deployment Servers....360
12.5.3 Delete a Bare Metal Cluster....361
12.6 Bare Metal Chassis....361
12.6.1 Add Bare Metal Chassis....361
12.6.1.1 Manually Add Bare Metal Chassis....361
12.6.1.2 Import Bare Metal Chassis from a Template....362

12.6.2 Bare Metal Chassis Status Management..... 363
12.6.3 Access a Bare Metal Chassis.....364
12.6.4 Delete a Bare Metal Chassis.....365
12.7 Bare Metal Instances..... 365
12.7.1 Create a Bare Metal Instance.....365
12.7.2 Bare Metal Instance Status Management.....367
12.7.3 Delete a Bare Metal Instance..... 368
Platform Configuration.....369
13.1 Overview.....369
13.2 Time Management..... 369
13.2.1 Modify Time Source Configuration..... 370
13.2.2 Synchronize Time..... 371
13.3 Log Server Management.....372
13.3.1 Add a Log Server..... 372
13.3.2 Log Server Configuration Options..... 373
13.3.3 Edit, Test, and Delete Log Servers..... 375
13.4 Email Server Management.....375
13.4.1 Add an Email Server..... 375
13.4.2 Edit, Enable or Disable, and Delete Email Servers..... 376
13.5 Theme and Appearance.....376
13.5.1 Customize Theme and Appearance..... 376
13.5.2 Restore Default Theme and Appearance..... 378
13.6 System Parameters.....378
13.6.1 Modify System Parameters.....378
13.6.2 Restore Default System Parameters..... 378
13.6.3 Parameter Entries..... 379
13.6.3.1 Platform Policy..... 379
13.6.3.2 Operations Management.....385
13.6.3.3 Host and VM..... 386
13.6.3.4 Image Storage.....389
13.6.3.5 Data Storage..... 390
13.6.3.6 Network Resources.....393
Glossary..... 394

1 Product Overview

1.1 What Is NexaSphere?

Through compute virtualization, storage virtualization, and network virtualization, NexaSphere organizes physical server, storage, and network resources into centrally managed data center resource pools, and provides lifecycle management capabilities for virtual machines, images, data storage, distributed networks, and related resources.

NexaSphere is designed for virtualization infrastructure deployment and VMware replacement scenarios. Administrators can use NexaSphere to build virtualized data centers, manage x86 and ARM servers, and manage virtual machines, hosts, data storage, networks, security, monitoring, backup, and bare metal capabilities in a unified interface.

This book helps you understand the role, architecture, resource model, and major capabilities of NexaSphere before you move on to installation, upgrade, resource management, operations, or security manuals.

Intended Audience

This book is intended for readers who need to quickly understand NexaSphere, including:

- Deployment and O&M engineers
- Technical support engineers
- Product consulting engineers
- Administrators who need to understand the resource model and functional scope of NexaSphere

1.2 NexaSphere Architecture and Resource Model

The architecture and resource model of NexaSphere explain how platform objects are organized, how they are associated with each other, and which resource boundaries administrators need to consider when planning and managing resources.

The architecture topic explains how NexaSphere organizes management nodes, data centers, clusters, hosts, virtual machines, storage, and networks into a virtualization platform. The resource model topic explains object hierarchy, attached resources, service dependencies, and scale relationships between these objects.

1.2.1 NexaSphere Architecture

NexaSphere provides the web console, resource orchestration, and platform management capabilities through management nodes. It organizes underlying compute resources through data centers, clusters, and hosts, and provides storage and network services required by virtual machines through data storage, image storage, and distributed networks.

Core Objects

- Management nodes run platform management services and provide a unified management entry and control plane.
- Data centers are the largest resource management domains in NexaSphere. They isolate and organize basic resources such as clusters, hosts, data storage, and distributed switches.
- Clusters are located under data centers and consist of hosts. Clusters run virtual machines and attach data storage and distributed switches.
- Hosts are physical servers that provide CPU, memory, local storage, and physical network resources. They provide the runtime environment for virtual machines.
- Virtual machines run on hosts and run business workloads through virtual CPUs, memory, disks, NICs, and peripheral devices.
- Data storage provides storage space for virtual machine disks, snapshots, and related data. It can come from local storage, NFS, SAN storage, or distributed storage.
- Distributed switches and distributed port groups provide unified virtual network configuration for virtual machine traffic and host-related traffic.
- Image storage and images provide the base image sources for virtual machine creation, template-based deployment, and batch delivery.

Runtime Relationship

A typical NexaSphere environment is organized around data centers, clusters, and hosts. Data storage and distributed switches are attached to clusters to provide storage and network services for virtual machines. A virtual machine is created from an image, stores disk data on data storage, and connects to virtual networks through distributed port groups.

1.2.2 NexaSphere Resource Model

The NexaSphere resource model describes object hierarchy, attached resources, service dependencies, and scale relationships. These relationships help administrators identify resource boundaries, understand dependency order, and prepare the required storage and network resources before virtual machines are deployed.

Table 1-1: Key Resource Relationships

Resource Object	Associated Object	Relationship Type	Description
Data Center	Cluster, Data Storage, Distributed Switch	Hierarchy	A data center defines a resource management domain. Basic resources in different data centers are isolated from each other.
Cluster	Host	Hierarchy	A cluster consists of hosts and is used to run and schedule virtual machine workloads.
Host	Virtual Machine	Runtime placement	Virtual machines run on hosts and consume CPU, memory, network, and other resources provided by hosts.
Cluster	Data Storage	Attachment	After data storage is attached to a cluster, it can provide disk data storage for virtual machines in the cluster.
Cluster	Distributed Switch	Attachment	After a distributed switch is attached to a cluster, it can provide unified virtual network configuration for virtual machines in the cluster.
Distributed Switch	Distributed Port Group	Hierarchy	A distributed port group is a logical group of distributed switch ports and carries network configurations such as VLAN, network segment, and DHCP.
Image Storage	Data Center, Data Storage	Service dependency	Image storage provides image services for data centers and works with data storage to support virtual machine creation and image usage.
VM Group	Virtual Machine	Grouping	VM groups are used to organize and manage virtual machines in a hierarchical structure within a data center.

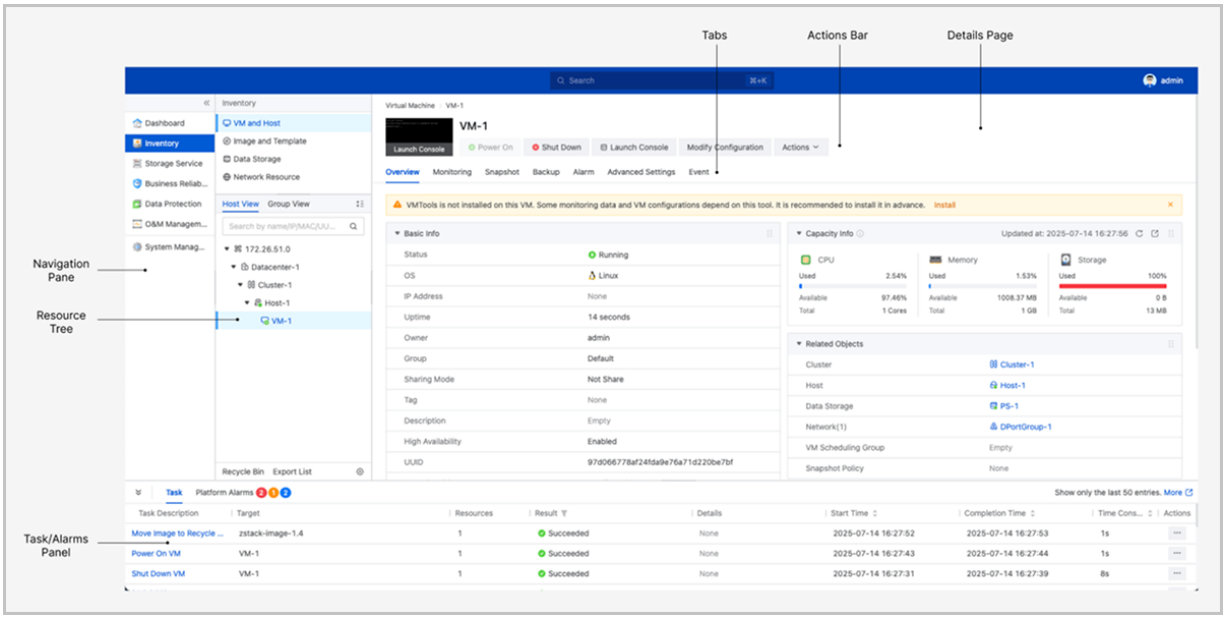
A data center can usually contain multiple clusters, data storage resources, and distributed switches. A cluster can contain multiple hosts and attach multiple related resources. A host can

run multiple virtual machines. Different storage types, network plans, and business scenarios affect how resources are combined and how many workloads they can support.

1.3 NexaSphere User Interface

NexaSphere provides a graphical management interface for viewing resources, configuring resources, and performing operations through the navigation pane, resource tree, details page, and task/alarms panel.

Figure 1-1: User Interface



- **Navigation Pane**: Located on the left of the interface, the navigation pane provides quick access to major platform functions, such as Dashboard, Inventory, Storage Service, and System Management. The navigation pane and resource tree form the primary navigation framework of the platform. You can switch between modules through the first-level and second-level menus in the navigation pane.
- **Resource Tree**: Located under the secondary menu of Inventory, the resource tree displays resources such as virtual machines, data storage, and distributed switches in a tree structure. You can browse and manage resources hierarchically, expand or collapse nodes, and use the right-click menu.
- **Task/Alarms Panel**: Located at the bottom of the interface, the task/alarms panel displays platform tasks and real-time platform alarms. It is globally visible so that you can track task execution and platform status changes. You can expand or collapse the panel.

- **Details Page**: The main workspace of the interface. The details page displays detailed information about the selected resource, including configuration overview, resource monitoring, and associated resources. It can be displayed as cards, lists, or other views.
- **Action Bar**: Located at the top or side of the details page, the action bar provides operation buttons for the selected resource. You can use the action bar to perform management actions.
- **Tabs**: Located near the top of the details page, tabs allow you to switch between different functional views on the same page. Related resource tabs and the action bar together support a select-view-operate workflow.

2 Installation and Initialization

2.1 Overview

About This Guide

This guide describes how to complete pre-deployment preparation, management node installation, management node HA setup, basic service checks, first login, and environment initialization for NexaSphere.

Intended Audience

This guide is intended for administrators, implementation engineers, and O&M engineers who are responsible for deploying and initializing NexaSphere. Readers should be familiar with server hardware, basic networking, operating system installation, and virtualization infrastructure deployment.

Installation Modes

NexaSphere provides three installation modes. Each includes different components and is intended for different scenarios.

Installation Mode	Installed Content	Applicable Scenario
Management Node Mode	• Base operating system • Dependencies such as `MariaDB` and `CloudBus` message bus	Suitable for installing a management node for the virtualization platform.
Compute Node Mode	• Base operating system • Virtualization components such as `libvirt` and `QEMU`	Suitable for installation as nodes other than the management node, for example: • Compute nodes • Standalone image storage, distributed image storage, and monitoring nodes of distributed image storage • Distributed storage nodes and monitoring nodes of distributed storage
Expert Mode	Base operating system	

2.2 Planning and Preparation

Before installing NexaSphere, complete hardware, network, software package, and server configuration preparation to ensure that the servers can boot from the installation media and meet the deployment requirements for management nodes, compute nodes, and subsequent business resources.

This chapter describes the basic conditions that must be confirmed before deployment, including hardware requirements, network planning, software package preparation, installation media creation, and server configuration. After completing these preparations, you can start installing the management node.

2.2.1 Hardware Requirements

Basic Hardware Requirements

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

- Basic demonstration environment: The server must have at least 4 CPU cores and 8 GB of memory.
- Production environment:
  - Management node configuration: In small-scale scenarios, the server must have at least 8 CPU cores, 16 GB of memory, and 240 GB of storage.

    The actual hardware configuration requirements depend on your business scale. For details, contact official technical support.
  - Compute node configuration: The actual hardware configuration requirements depend on your business scale. For details, contact official technical support.

Recommended Hardware Configuration

The following table lists the recommended server hardware configuration for both basic demonstration environments and production environments:

Device		Configuration
Server	CPU	• x86 architecture: 64-bit. Supports Intel VT-x or AMD-V hardware virtualization features (such as Intel VMX or AMD/Hygon SVM). • ARM architecture: 64-bit. Supports hardware virtualization features.

Device		Configuration
	Memory	No special requirements. DDR4 or higher specification memory is recommended.
	Motherboard	Standard dual-socket server motherboard.
	RAID Controller	Supports SAS/SATA RAID 0/1/10 and passthrough mode .
	Disk	No special requirements. Select HDD or SSD based on capacity and performance requirements.
	Network Ports	• Management port: 1 × 1 GbE RJ45 • Business port: 1 × 10 GbE SFP+
Network Switch		• At least one 1 GbE switch. 10 GbE switch is recommended. • Multiple Category 5 or higher network cables.

Hardware Planning

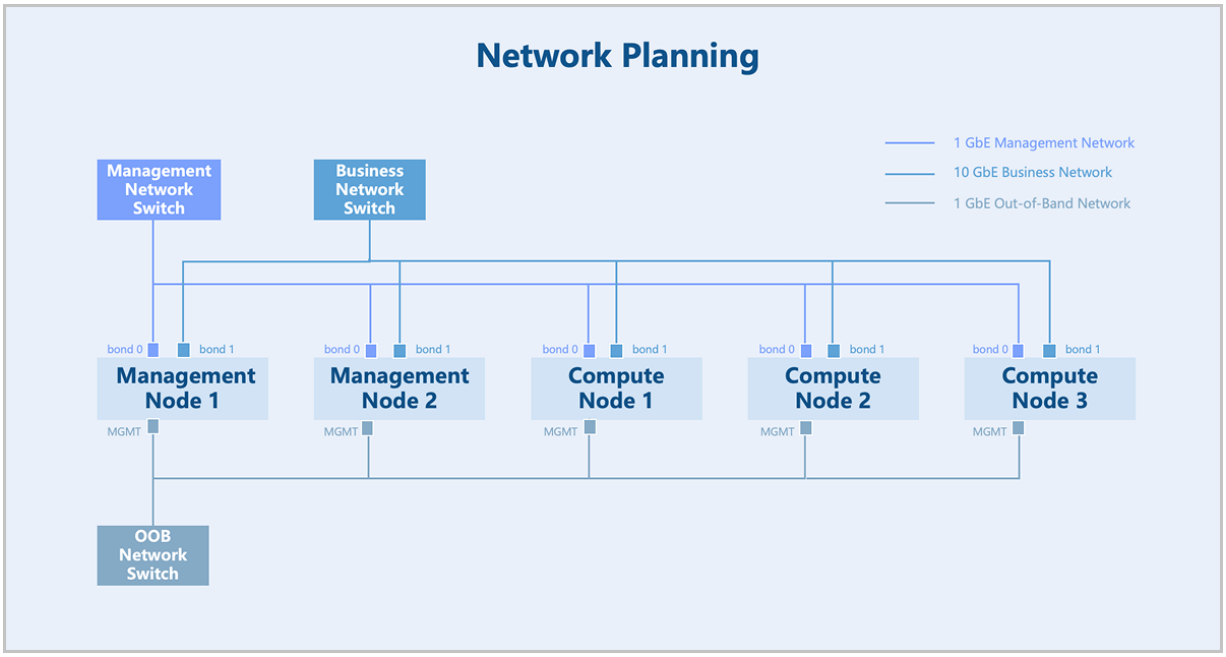
You need to plan your server resources according to your production needs. In a large data center with sufficient resources, we recommend deploying two servers as management nodes to control the entire virtualization platform. These two management nodes provide high availability. If one management node fails, the system automatically triggers a high availability switchover within seconds. This ensures the management service remains continuously available. In a small data center, you can use a single server as the management node.

The remaining servers act as compute nodes. In a large data center, you can select multiple compute nodes to act as image storage. This expands the total image storage capacity and improves throughput. In a small data center, the image storage can share a server with the management node.

2.2.2 Network Planning

A typical data center consists of several servers with different functions and two independent network environments.

Figure 2-1: Network Planning



- The management network uses dual Gigabit networks for managing hardware resources related to the virtualization platform. The business network uses dual 10-Gigabit networks for virtual machine service traffic. You can configure network speeds based on your actual business requirements.
- If you use network storage such as NFS or distributed storage, you need also plan an additional storage network to handle storage traffic.
- We recommend that you maintain consistent NIC naming across all servers and use NICs with the same name to carry the same type of network traffic. For example, all management network traffic use the `em1` NICs.
- If you use a VLAN network, you need to configure the corresponding VLAN network communication on the switch in advance.
- NexaSphere automatically assigns IP addresses to virtual machines. You need to reserve an IP range that does not conflict with the system and ensure this range does not conflict with any existing DHCP services in your network environment.

2.2.3 Prepare Software Packages

Software Package Description

NexaSphere ISO is built on Helix, a self-developed virtualization kernel that runs between the infrastructure layer and upper-layer guest operating systems. Helix integrates hardware drivers,

macro kernel components, virtual agents, and other required components. When guest operating systems access underlying hardware resources, Helix masks differences across heterogeneous hardware, removes the dependency on physical hardware drivers, and improves hardware compatibility, reliability, availability, scalability, and performance in the virtualization environment.

Installing NexaSphere does not require Internet access or a yum repository. The installation can be completed offline.

Software Package List

The following table lists the software packages used to install NexaSphere:

Server Architecture	Software Package Name
x86	NexaSphere-x86_64-DVD-2.1.1-H84r.iso
	NexaSphere-x86_64-DVD-2.1.1-KylinV10SP3.iso
ARM	NexaSphere-aarch64-DVD-2.1.1-H22e.iso
	NexaSphere-aarch64-DVD-2.1.1-KylinV10SP3.iso

Note:
After obtaining the software, use an MD5 checksum tool to verify the checksum and ensure that the software is complete and intact.

2.2.4 Create Installation Media

After obtaining the NexaSphere ISO image file, write the ISO image to a USB drive so that the physical server can boot the installer from the USB drive.

You can create the installation media with Rufus or Fedora Media Writer based on your local operating system and tool preference. After the installation media is created, set the USB drive as the boot media in the server BIOS or boot menu.

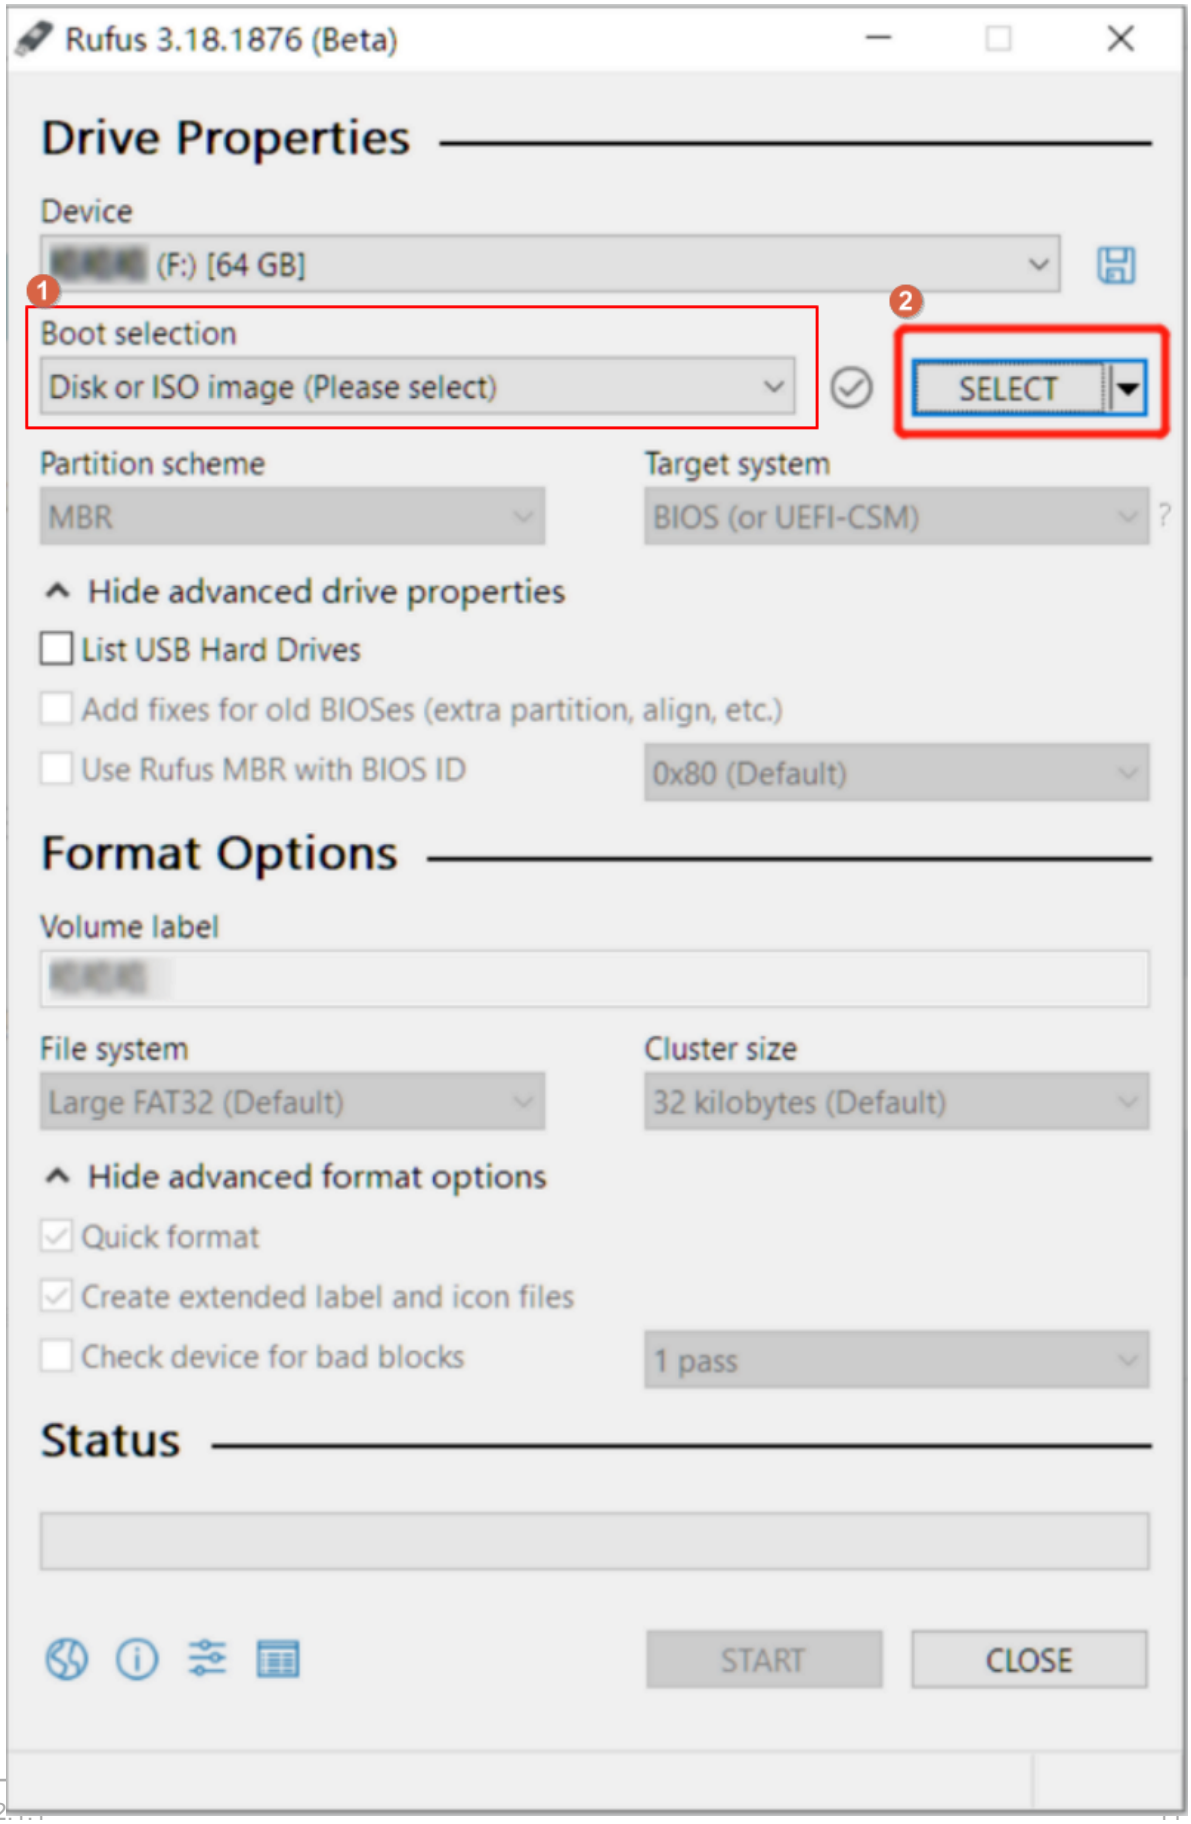
2.2.4.1 Create a USB Drive with Rufus

After you obtain the ISO package, you can use Rufus to burn the ISO image to a USB drive.

Procedure

1. Select an ISO image.
   a) Connect your USB drive and open Rufus.
   b) In the **Boot Selection** drop-down list, choose **Disk or ISO image**.
   c) Click **SELECT** to open the ISO image file that you obtained.

Figure 2-2: Select ISO Image



2. Burn the image.

   a) After selecting the ISO image, keep the other options at their default settings and click **START**.

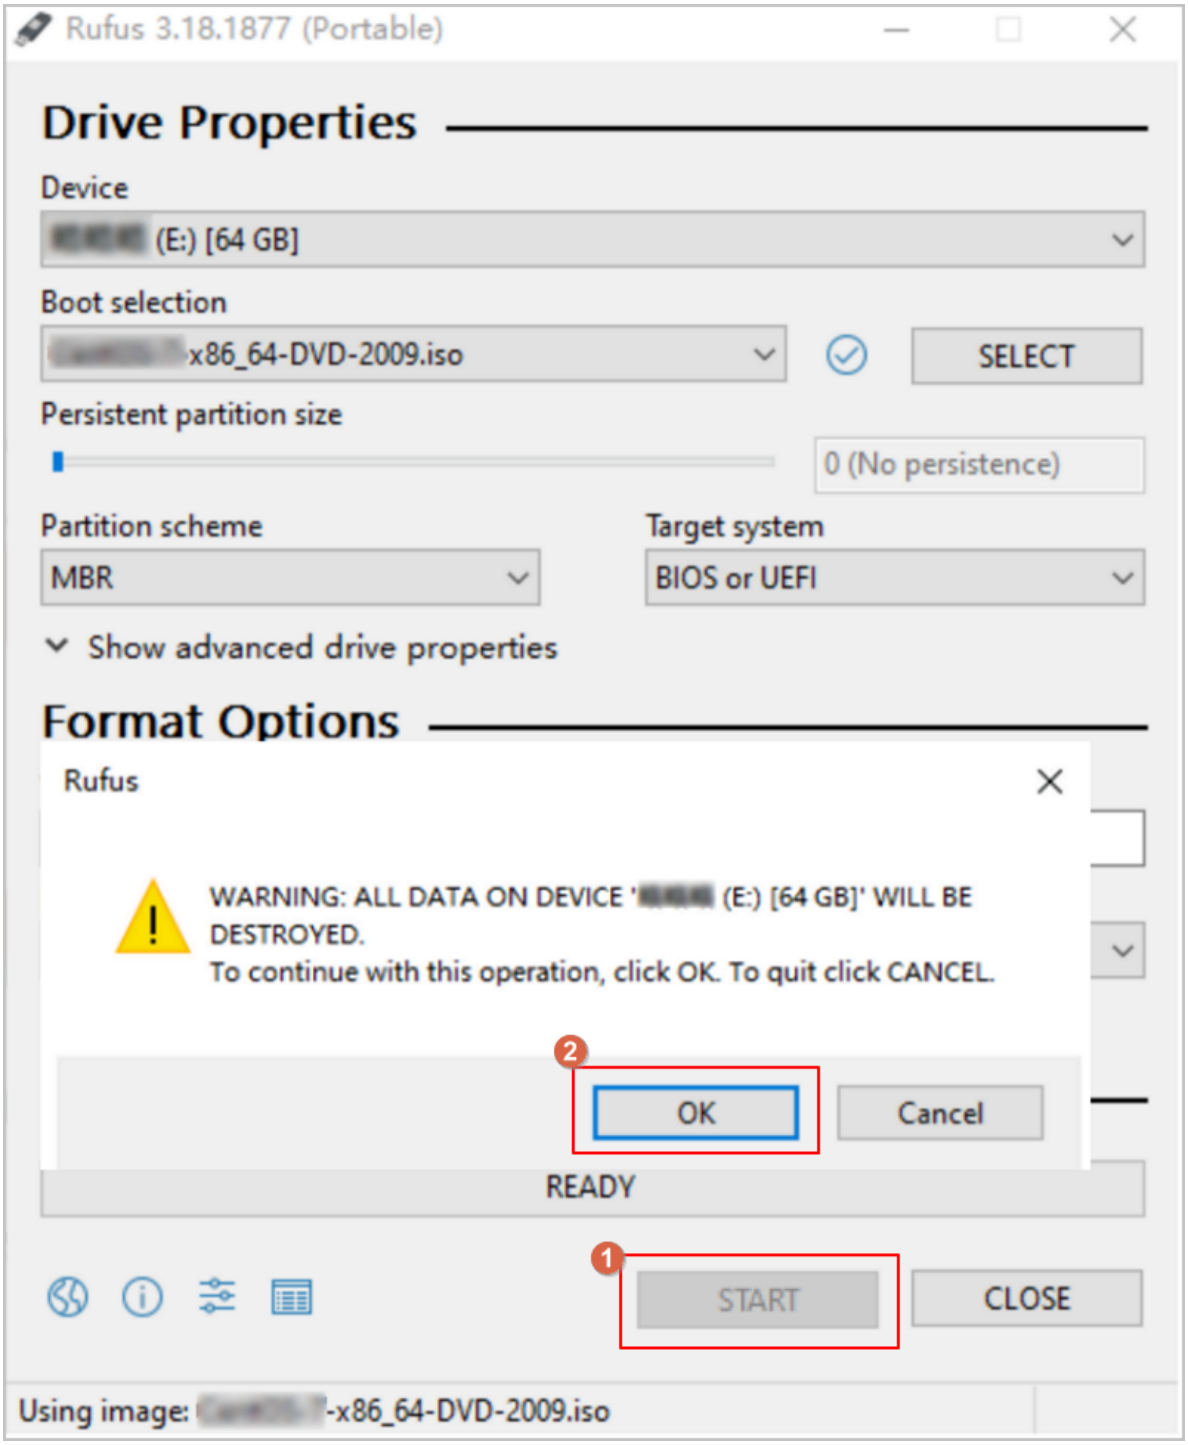
   b) Confirm the warning message.

   Note:

   Burning the image will format the data on the USB drive. If you have important data on your USB drive, make sure to back it up before formatting.

   c) Click **OK**, and Rufus will burn the ISO image to the USB drive.

Figure 2-3: Burn the Image



What's next

After the burning process completes, you can use the USB drive as a boot disk. NexaSphere supports booting in Legacy mode or UEFI mode.

2.2.4.2 Create a USB Drive with Fedora Media Writer

After you obtain the ISO package, if you are using the Kylin operating system, we recommend using Fedora Media Writer to burn the ISO image.

Procedure

1. Select an ISO image.
   a) Connect your USB drive and open Fedora Media Writer.
   b) Click **Custom Image** to open the ISO image file that you obtained.
2. Burn the image.
   a) After selecting the ISO image, keep the other options at their default settings and click **Write Disk**.

> **Note:**
> Burning the image will format the data on the USB drive. If you have important data on your USB drive, make sure to back it up before formatting.

What's next

After the burning process completes, you can use the USB drive as a boot disk. NexaSphere supports booting in Legacy mode or UEFI mode.

2.2.5 Configure Servers

Server Configuration Check

Based on your actual network planning, rack, cable, and power on the servers and network equipment. Then, start the servers, enter BIOS, and verify the following:

- The installation process overwrites all data, so confirm that you have backed up all disks in the servers.
- Enter the server BIOS and enable CPU VT and Hyper-Threading (HT) options.
- Configure the appropriate RAID level in the RAID controller to provide data redundancy.
- Set the USB drive as the first boot device.

For information about how to change server configurations, refer to your hardware vendor documentation.

2.3 Install the Management Node

The management node hosts NexaSphere management services, Web UI, and basic control services. It is the entry point for completing platform deployment and managing resources after deployment.

Select an installation method based on your network planning and deployment scenario. In production environments, bonded NIC deployment is typically used to improve management network reliability. In simple verification scenarios or single-NIC environments, you can use single NIC deployment. If you need to configure the management node from a text interface after installation, use the TUI installation path.

After installation, check the NexaSphere service status and determine whether to set up management node HA based on business continuity requirements.

2.3.1 Bonded NIC Deployment

To meet the requirements for network port bandwidth capacity and high reliability, NIC bonding must be configured in production environments. Since the operating system installation process is identical for both x86 and ARM servers, this chapter uses x86 servers to introduce the detailed installation procedure.

Procedure

1. Select the boot option.

   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.

   - GUI method
   - VNC method
   - Text mode method

2. Review the installation configuration summary.

   This page displays the system installation configuration. You can modify the configuration as needed. By default, NexaSphere is configured with the following settings:

   - **Keyboard**: English (US)
   - **Language Support**: English (United States)

- **Time & Date**: Asia/Shanghai (UTC+8). We recommend that you check the host's time in advance and configure it to the current time and time zone.

3. Select the installation mode.
   a) On the **INSTALLATION SUMMARY** page, click **Software Selection**.
   b) On the **SOFTWARE SELECTION** page, choose the installation mode.

      Set the servers for management nodes to **Management Node** mode. Set the other servers to **Compute Node** mode.

   c) After selecting the installation mode, click **Done**.
4. Configure the disk partitions.
   a) On the **INSTALLATION SUMMARY** page, click **Installation Destination** to enter the **INSTALLATION DESTINATION** page.
   b) For **Device Selection**, we recommend that you only configure the system disk. After the system is installed, you can configure other disks.
   c) 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:

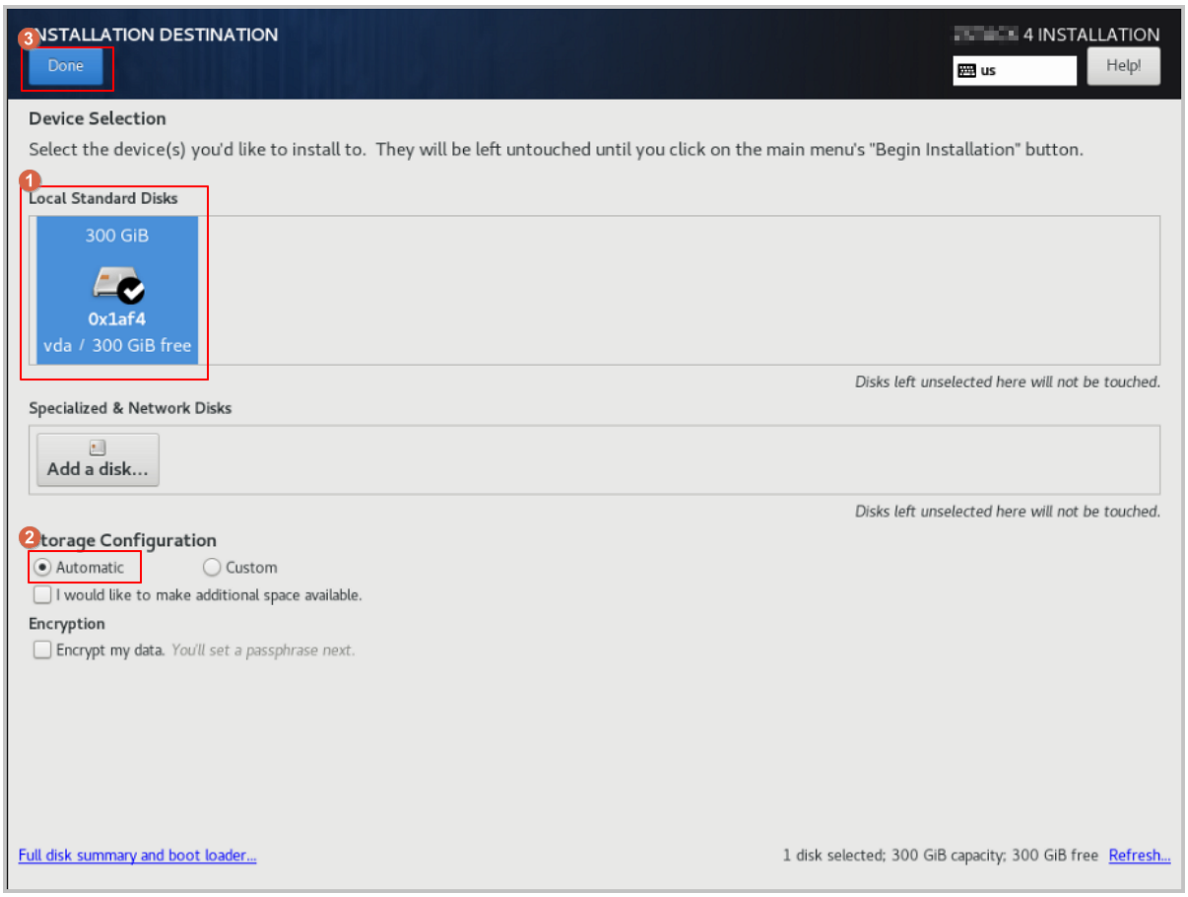
      - UEFI Mode:
        - `/boot`: This directory stores the core files needed for Linux boot. We recommend allocating 1 GB of space.
        - `/boot/efi`: This directory stores the UEFI boot files. We recommend allocating 500 MB.
        - *swap*: This is the swap area. We recommend allocating 32 GB.
        - `/`: 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 1 GB of space.
        - *swap*: This is the swap area. We recommend allocating 32 GB.
        - `/`: 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 NexaSphere (total disk capacity should be greater than 300 GB).
- In Legacy mode, if the system disk capacity exceeds 2 TB, you need to configure a BIOS boot partition to support GPT partitioning. UEFI mode does not have this limitation and supports GPT partitioning.

d) Review the configuration and click **Done**.

Figure 2-4: Configure Disk Partitions

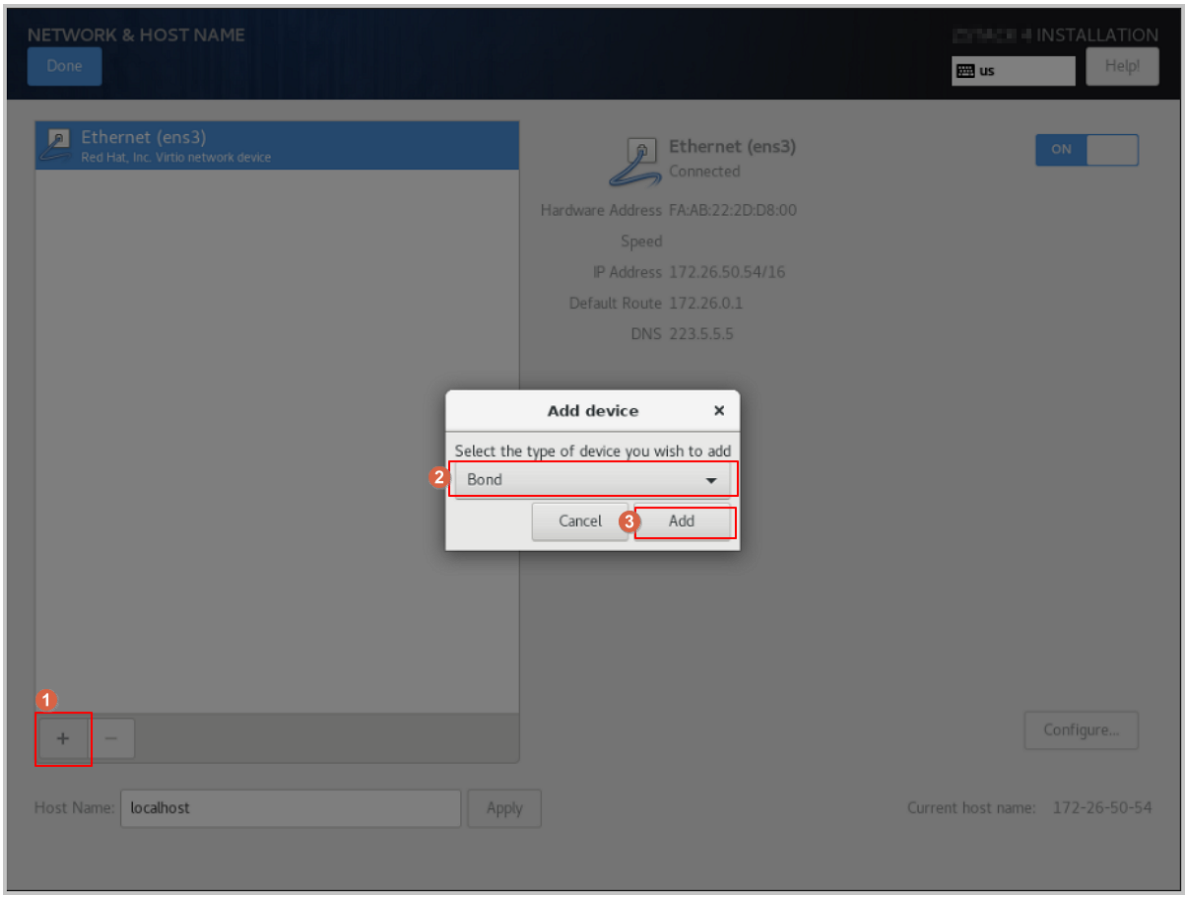


5. Add a bond device.

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 2-5: Add Bond Device

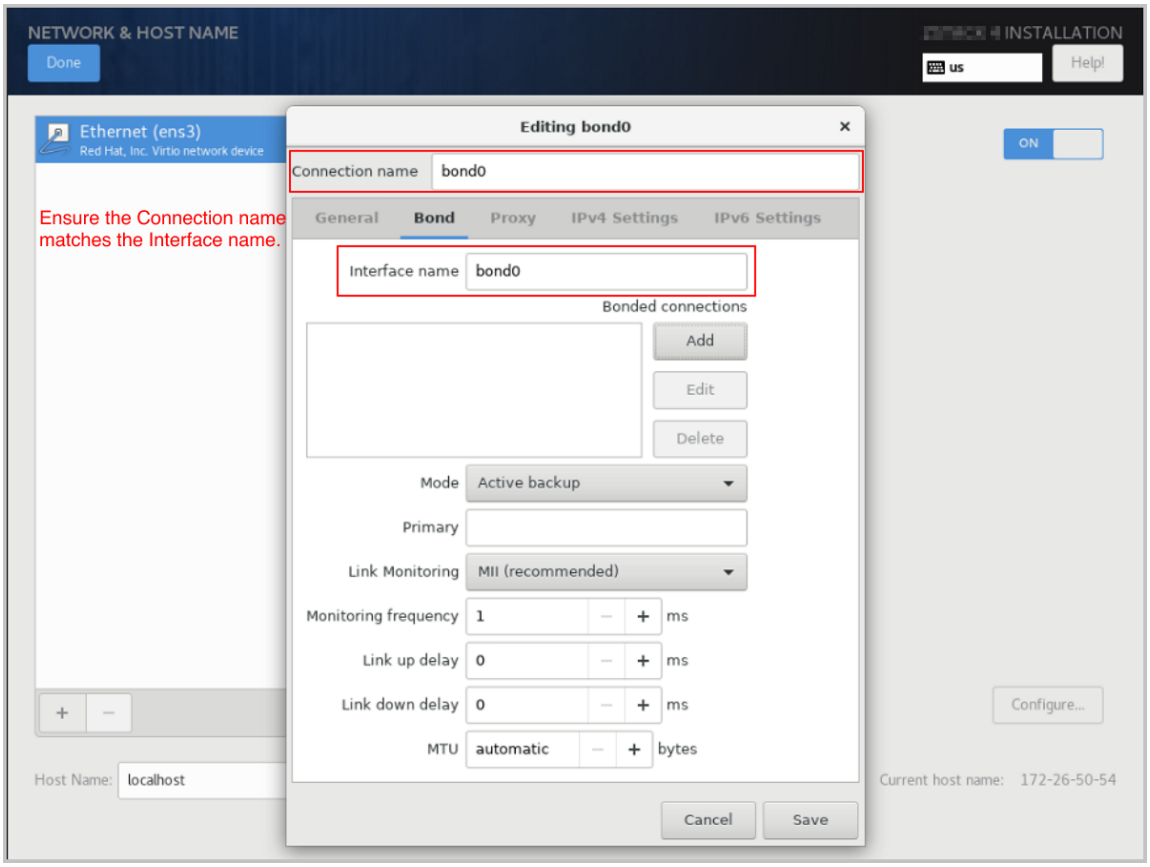


6. Add a Bond Slave.

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

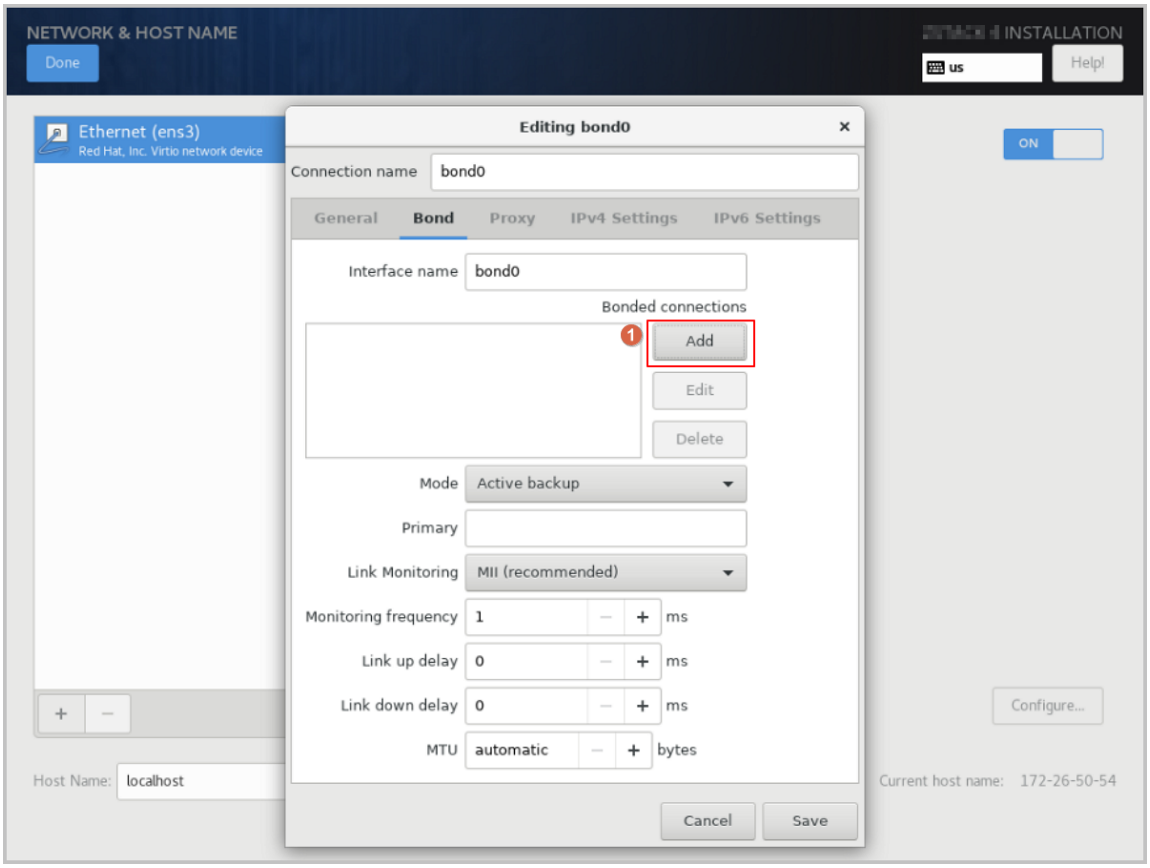
> **Note:**
>
> **Make sure the Connection name matches the Interface name.**

Figure 2-6: Adjust Connection Name



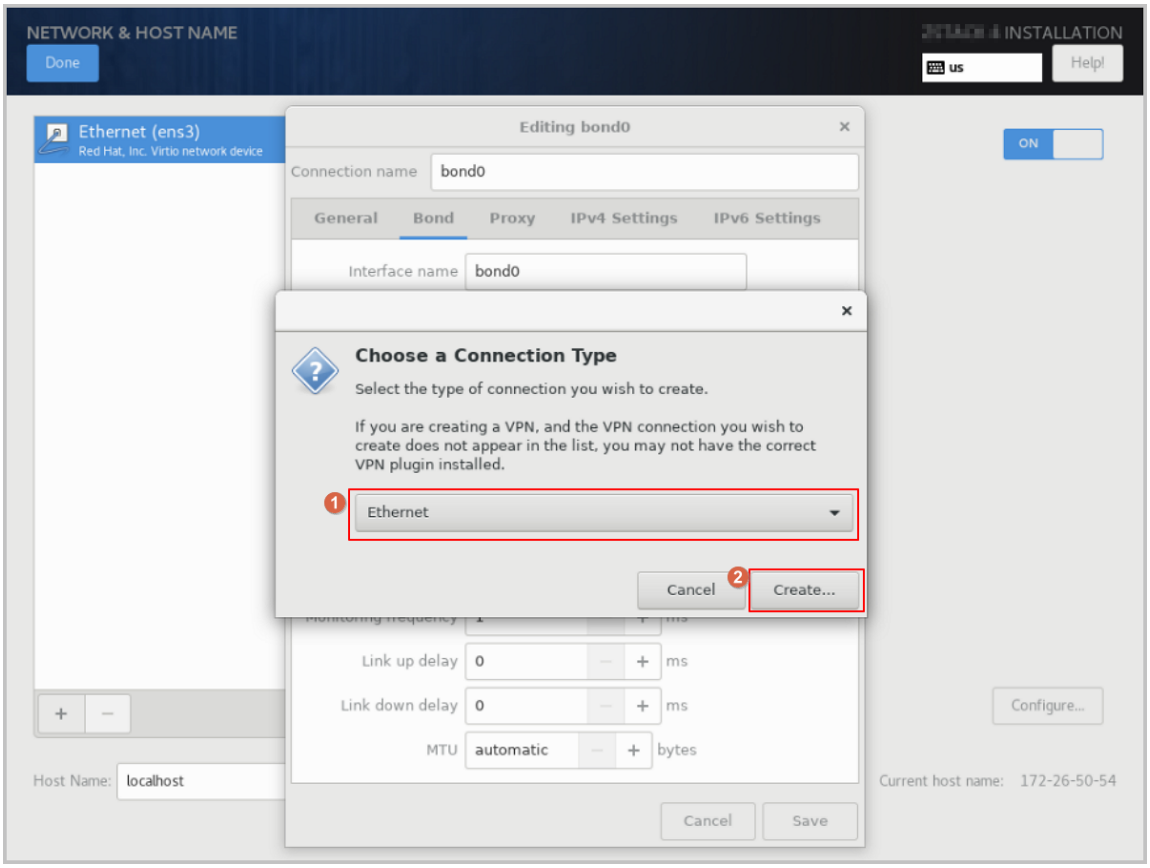
b) On the Bond configuration window, click **Add** to add a bond slave.

Figure 2-7: Add Bond Slave



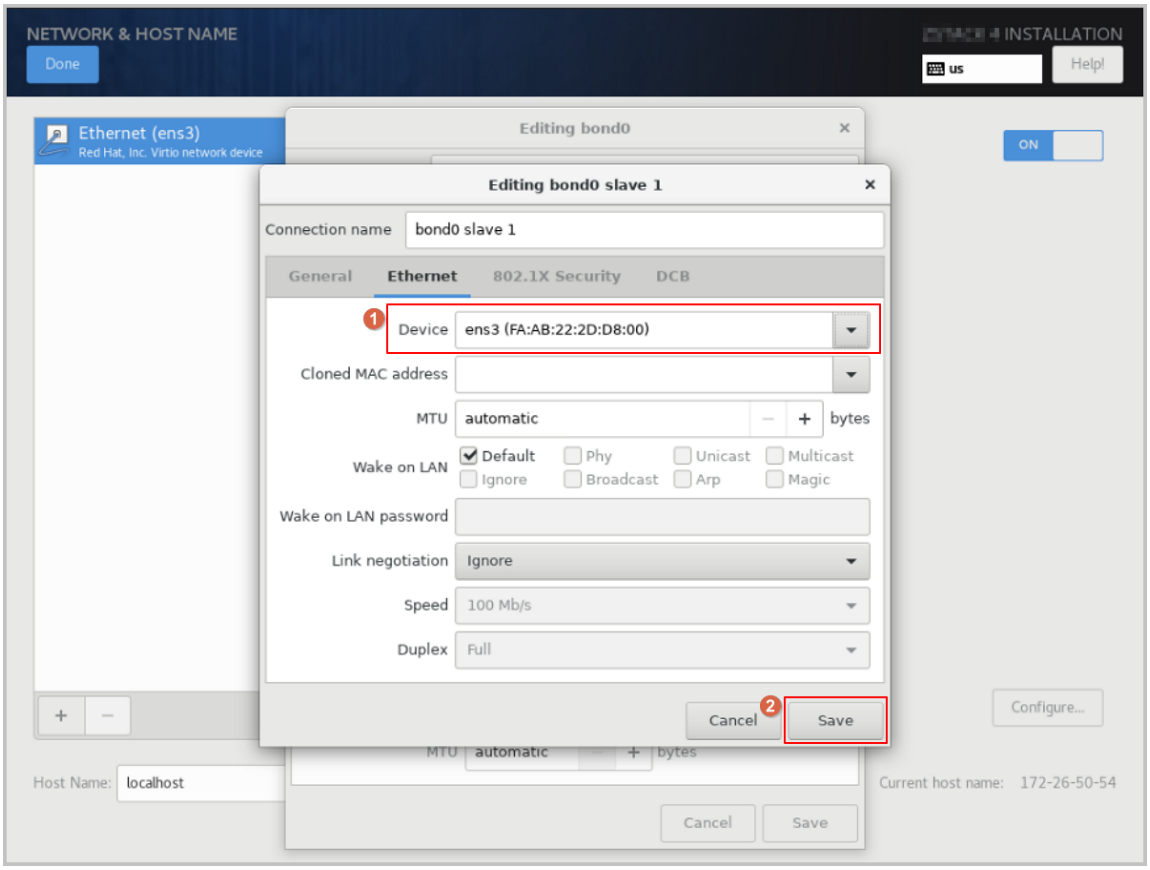
c) In the **Choose a Connection Type** dialog, choose a connection type from the drop-down list, such as **Ethernet**, and then click **Create...**.

Figure 2-8: 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**.

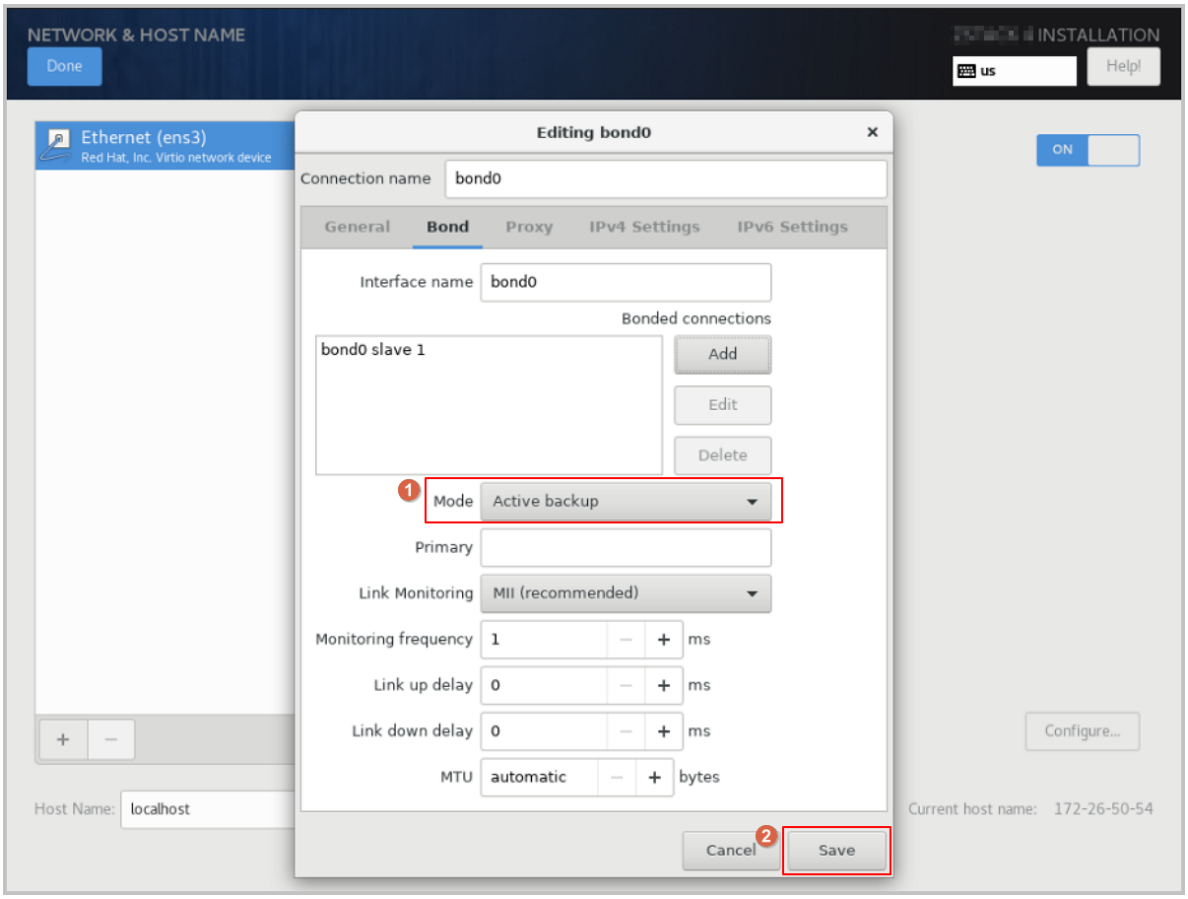
Figure 2-9: Select Bond Slave Device



7. 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**.

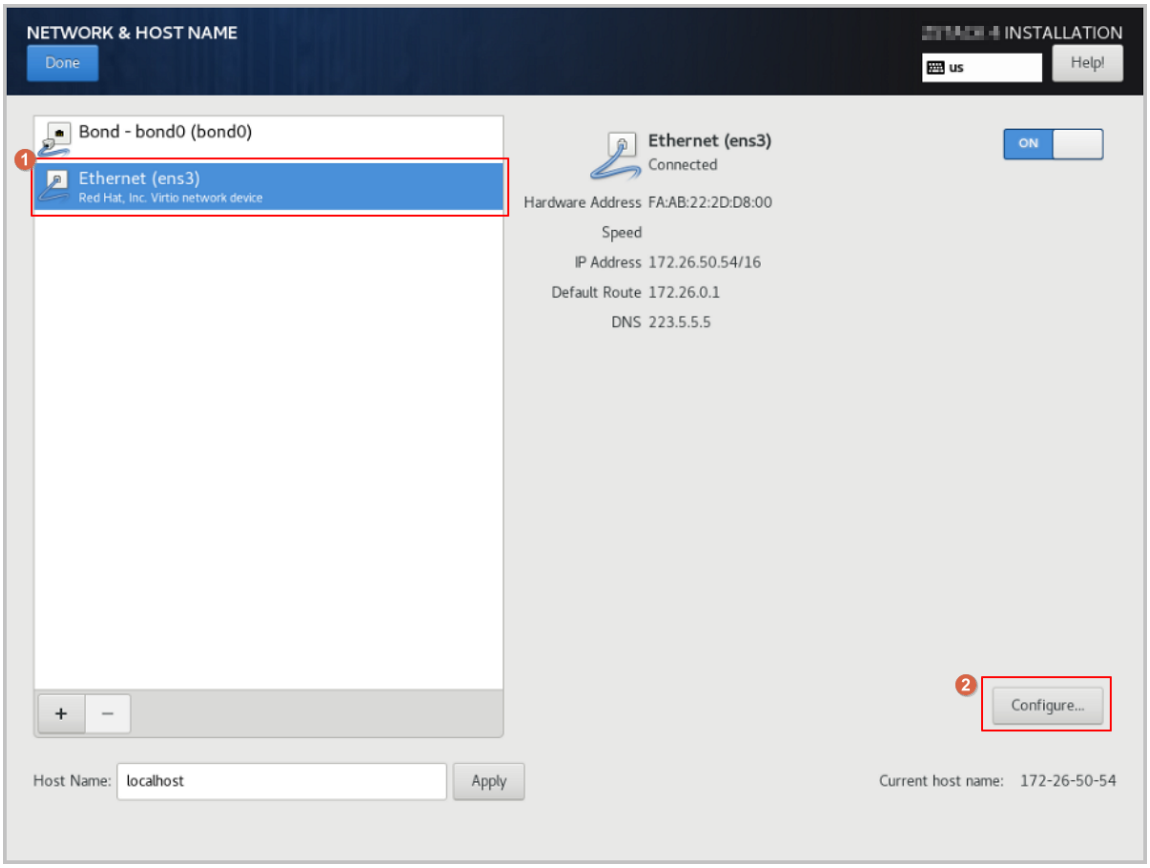
Figure 2-10: Select Bond Mode



8. Disable IPv4 on the original NIC.

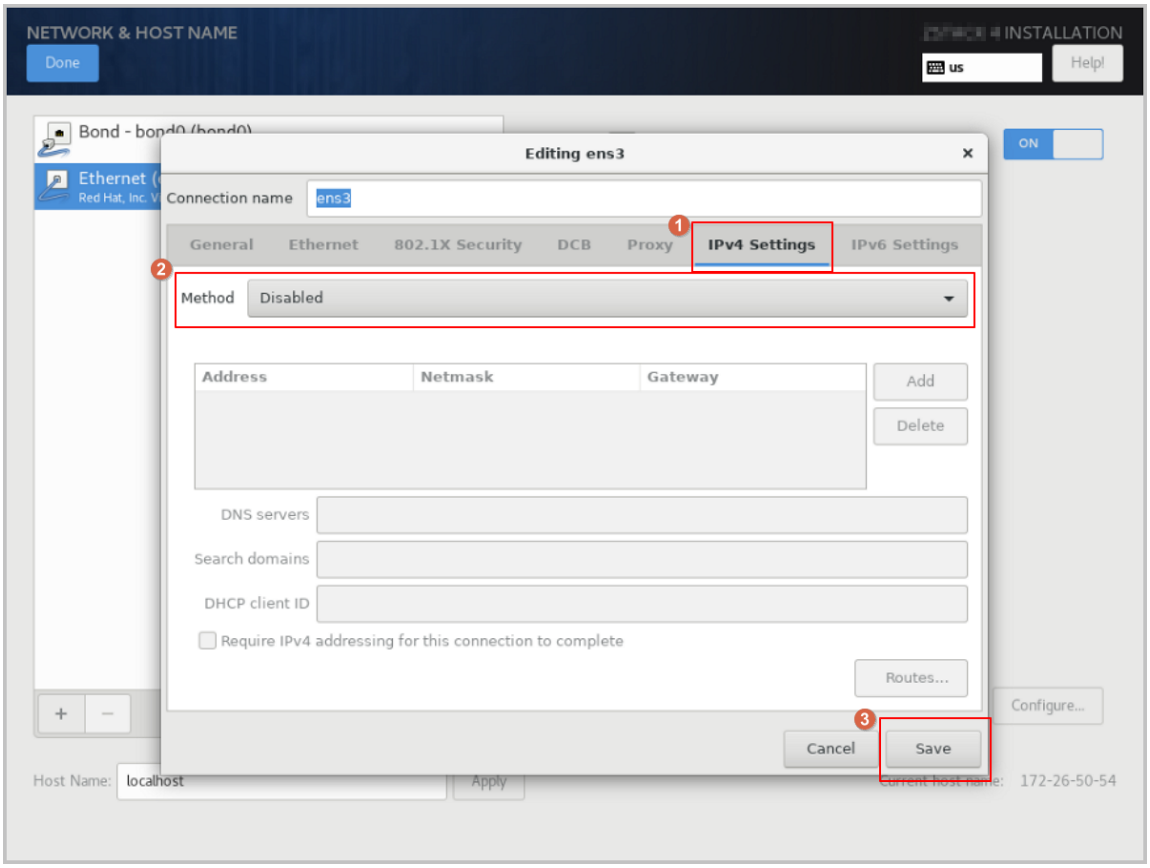
   a) Select the original NIC, such as **Ethernet (ens3)**, and click **Configure**.

Figure 2-11: 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 2-12: Disable IPv4



9. Configure a static IP address for Bond.

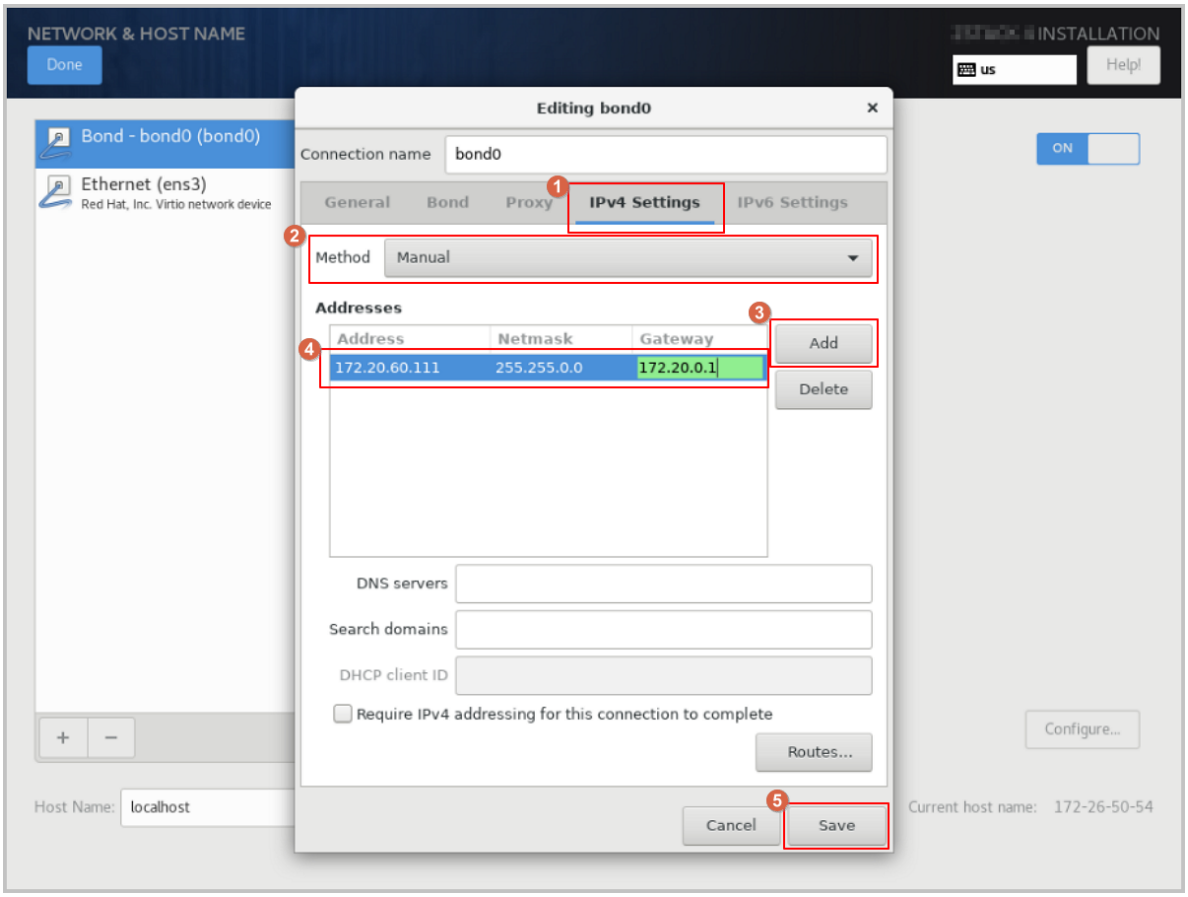
   a) On the **NETWORK & HOST NAME** page, choose the bond device, such as **Bond (bond0)**, and then click **Configure**.

   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.

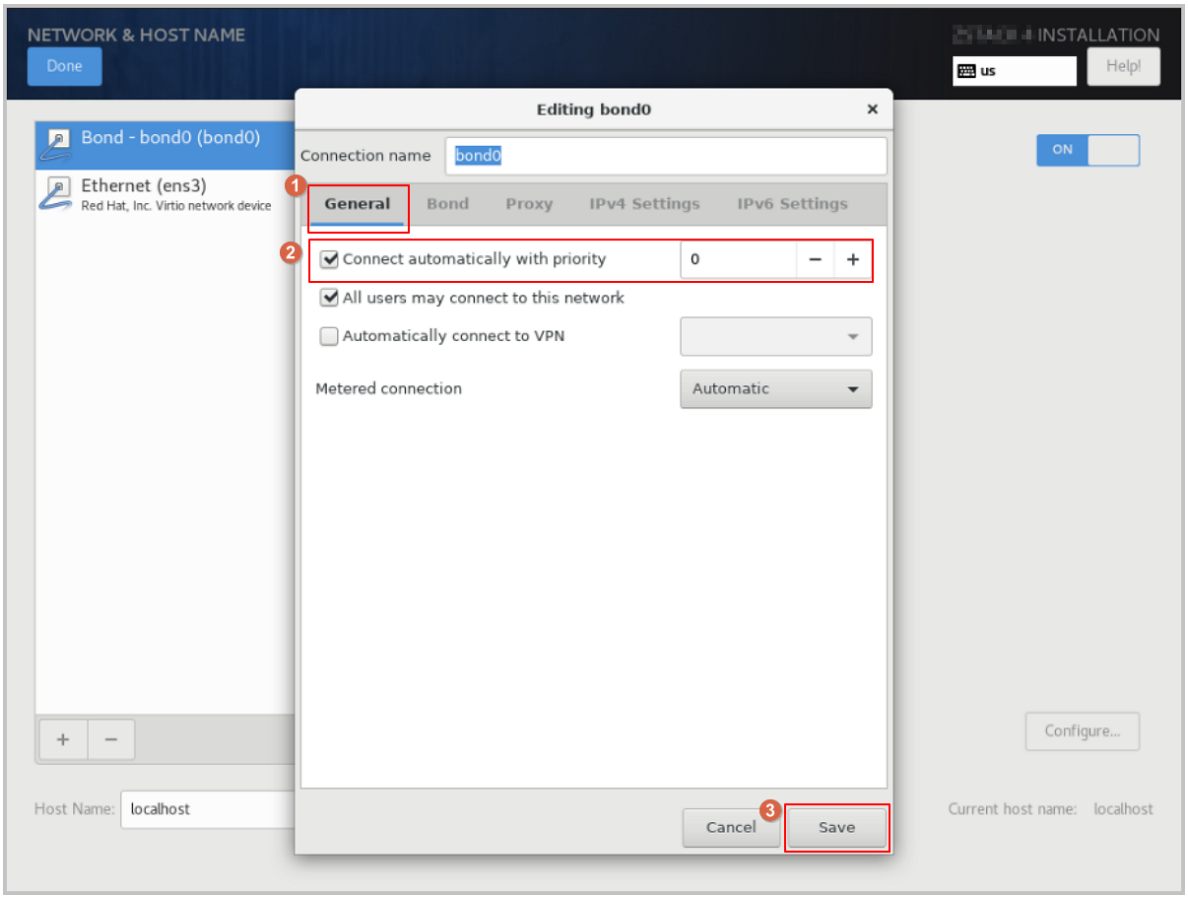
Figure 2-13: Configure Bond Static IP Address



10. 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 2-14: Configure NIC to Activate Automatically

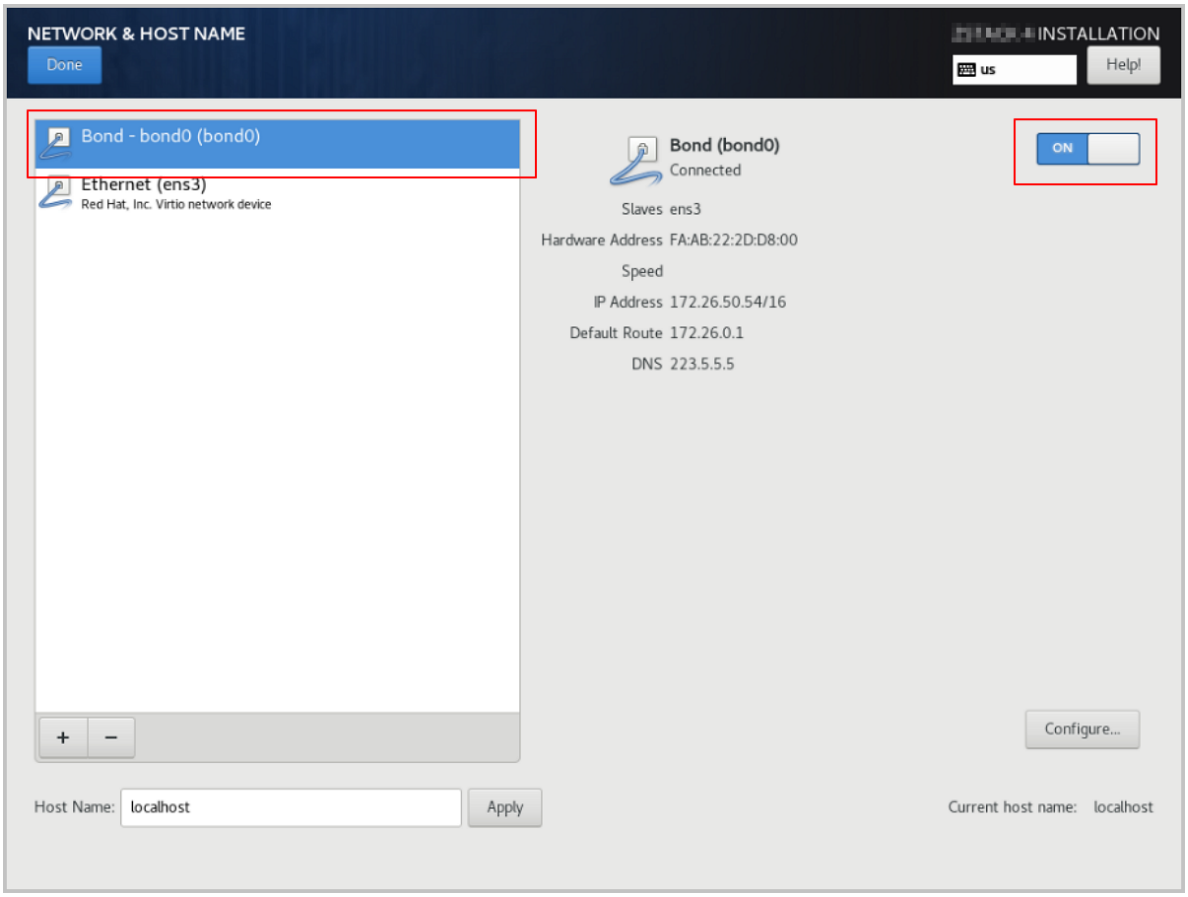


11. Complete the Bond configuration.

a) 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, NexaSphere will not be installed properly.

b) Review the configuration and click **Done** to return to the **INSTALLATION SUMMARY** page.

Figure 2-15: Check Bond Configuration



12. On the **INSTALLATION SUMMARY** page, click **Root Password** to set the root password for the operating system.
13. On the **INSTALLATION SUMMARY** page, click **Begin Installation** to begin installing the operating system.

2.3.2 Single NIC Deployment

The operating system installation process is identical for both x86 and ARM servers. This chapter uses x86 servers to introduce the detailed installation procedure.

Procedure

1. Select the boot option.

   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.

- GUI method
- VNC method
- Text mode method

2. Review the installation configuration summary.

   This page displays the system installation configuration. You can modify the configuration as needed. By default, NexaSphere is configured with the following settings:

   - **Keyboard**: English (US)
   - **Language Support**: English (United States)
   - **Time & Date**: Asia/Shanghai (UTC+8). We recommend that you check the host's time in advance and configure it to the current time and time zone.

3. Select the installation mode.

   a) On the **INSTALLATION SUMMARY** page, click **Software Selection**.

   b) On the **SOFTWARE SELECTION** page, choose the installation mode.

      Set the servers for management nodes to **Management Node** mode. Set the other servers to **Compute Node** mode.

   c) After selecting the installation mode, click **Done**.

4. Configure the disk partitions.

   a) On the **INSTALLATION SUMMARY** page, click **Installation Destination** to enter the **INSTALLATION DESTINATION** page.

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

   c) 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 1 GB of space.
        - `/boot/efi`: This directory stores the UEFI boot files. We recommend allocating 500 MB.
        - *swap*: This is the swap area. We recommend allocating 32 GB.

- /: This is the root directory for the Linux system. We recommend allocating all remaining space.

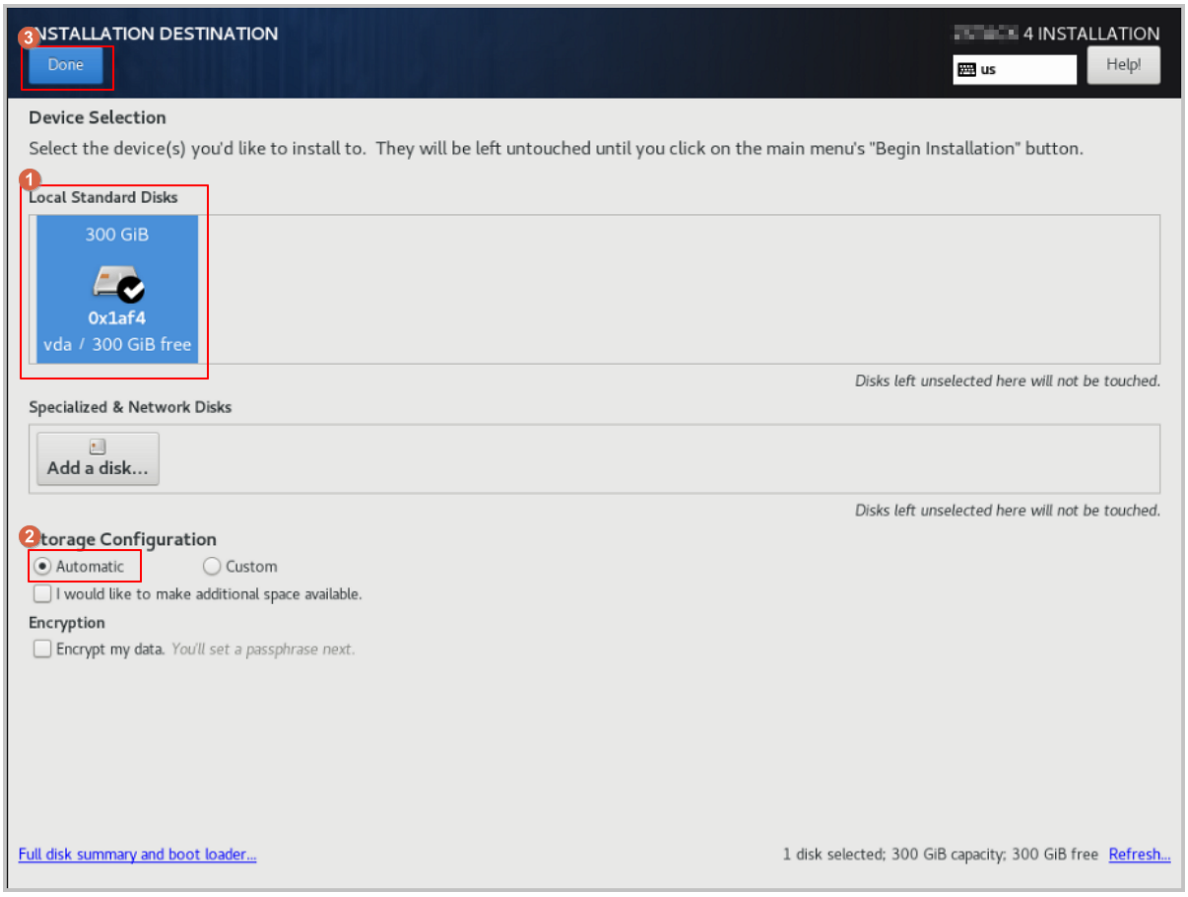
- Legacy Mode:
  - /boot: This directory stores the core files needed for Linux boot. We recommend allocating 1 GB of space.
  - swap: This is the swap area. We recommend allocating 32 GB.
  - /: This is the root directory for the Linux system. We recommend allocating all remaining space.

Note:

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

d) Review the configuration and click **Done**.

Figure 2-16: Configure Disk Partitions



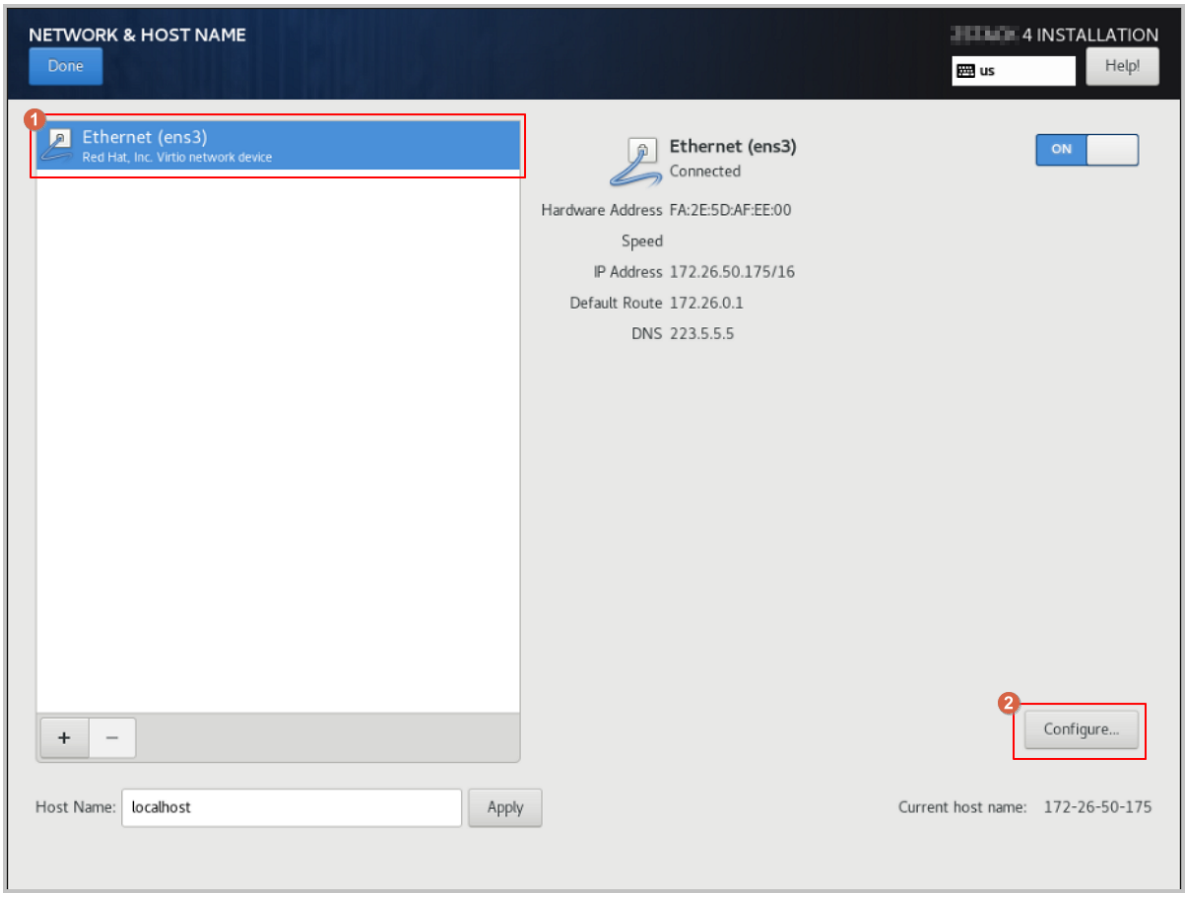
5. Start configuring the NIC.

   a) On the **INSTALLATION SUMMARY** page, click **Network & Host Name** to access the **NETWORK & HOST NAME** page.

   b) Select a NIC from the list on the left, such as **Ethernet (ens3)**.

   c) Click **Configure**.

Figure 2-17: Configure NIC



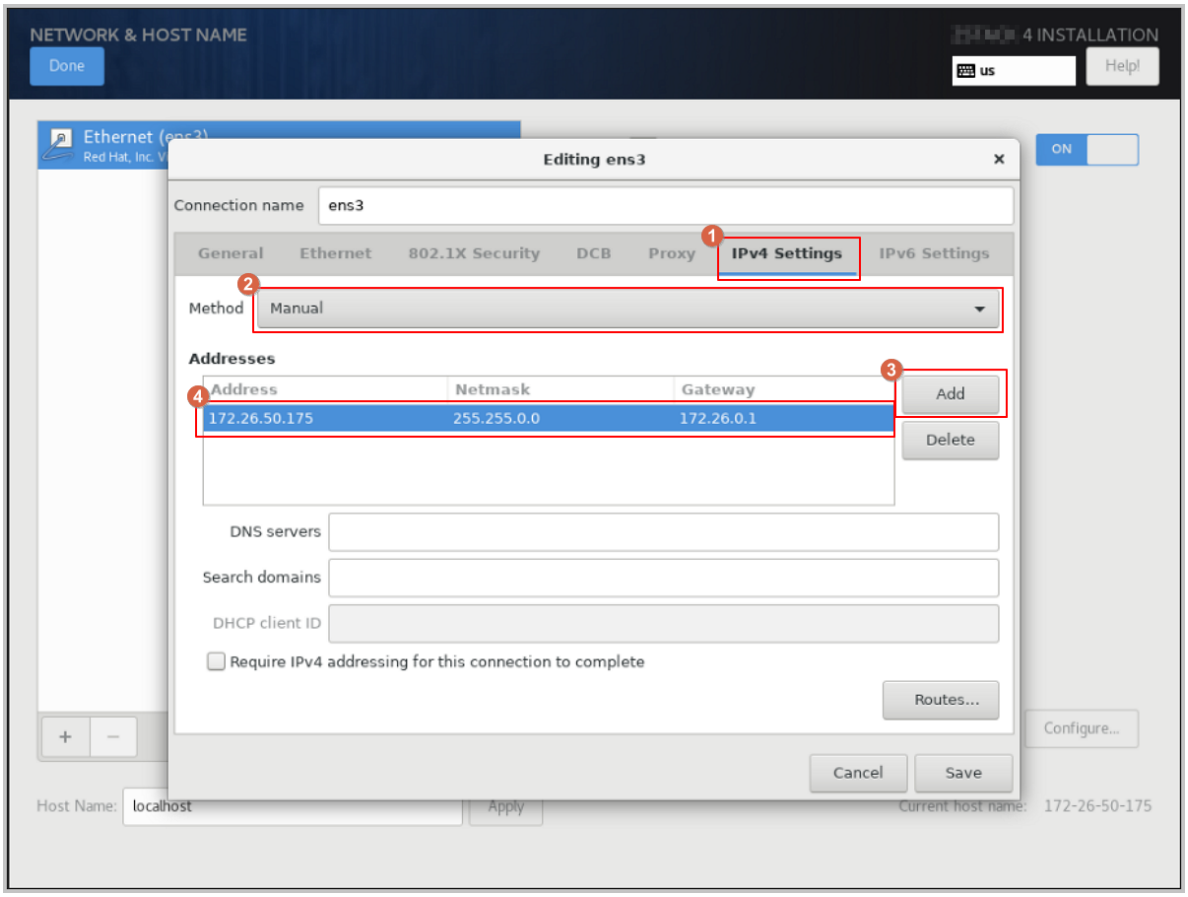
6. Configure a static IP address for the NIC.

   a) In the **Editing ens3** dialog, click **IPv4 Settings**.

   b) For **Method**, choose the IP address acquisition method as needed. For example, select **Manual** to specify the IP address manually.

   c) Click **Add** to add an IP address entry, and configure the IP address, netmask, and gateway as needed.

Figure 2-18: Configure Static IP Address

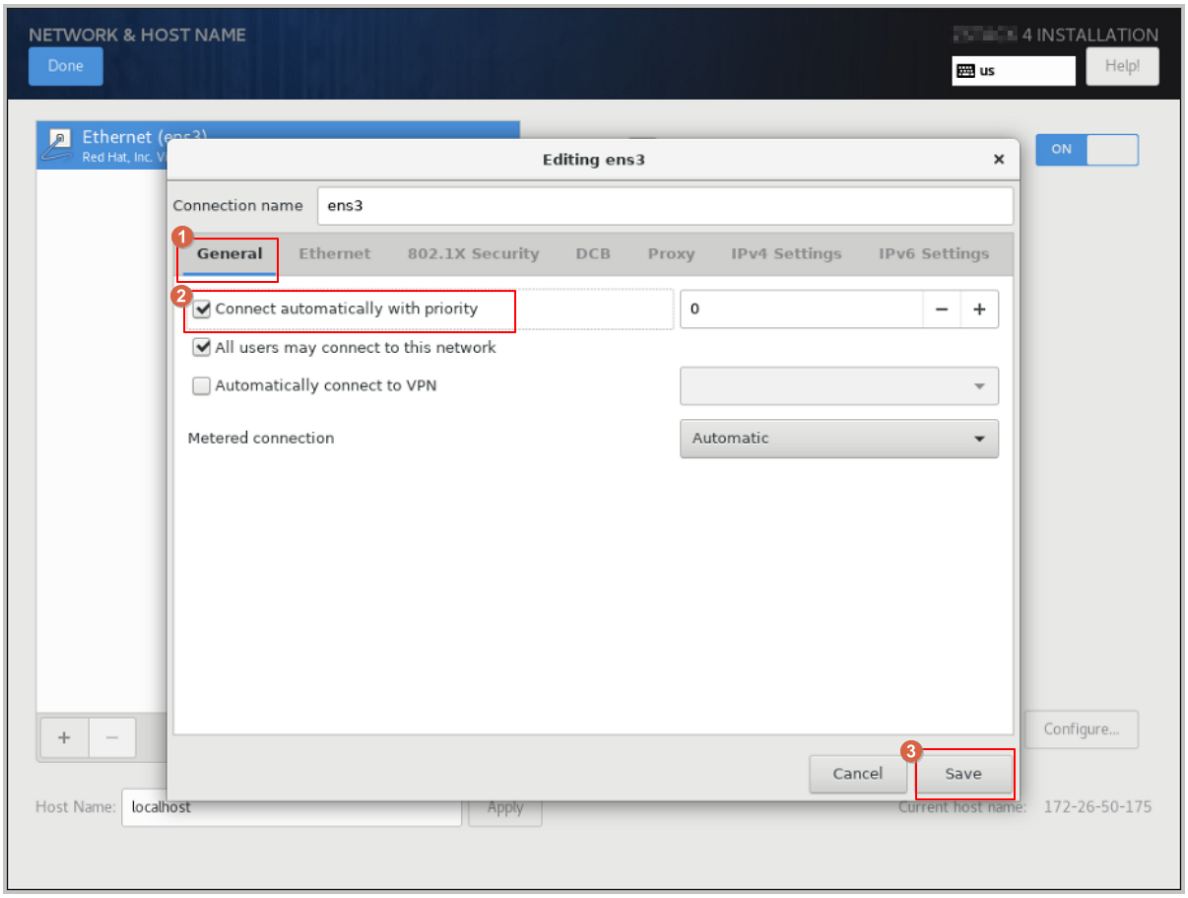


7. Configure the NIC to activate automatically.

   a) In the **Editing ens3** dialog, click **General**, then select the **Connect automatically with priority** checkbox to set the NIC for automatic activation.

   b) Review the configuration and click **Save**.

Figure 2-19: Configure NIC to Activate Automatically

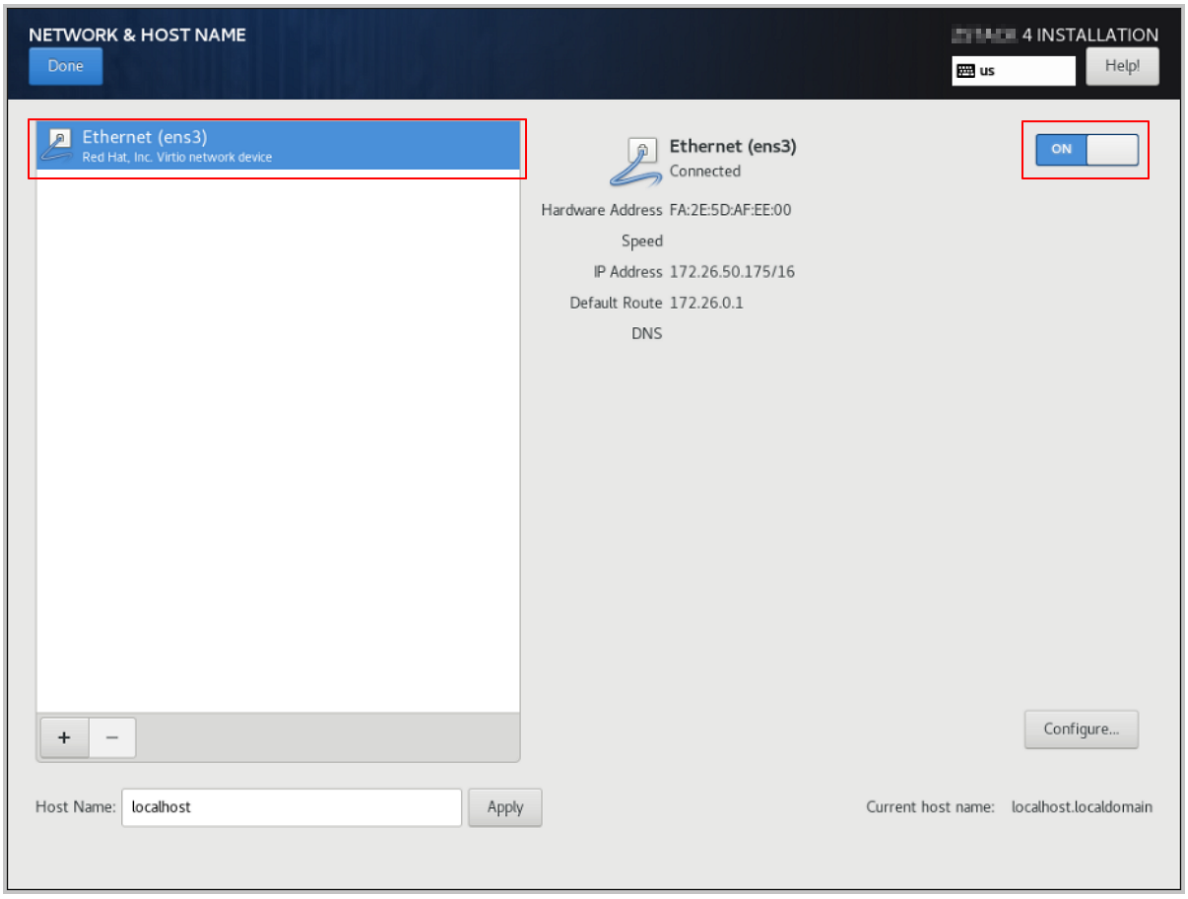


8. Complete the NIC configuration.

   a) Return to the **NETWORK & HOST NAME** page, and confirm that you have selected the correct NIC and that the NIC status is **ON**.

   b) Click **Done** to return the **INSTALLATION SUMMARY** page.

Figure 2-20: Check NIC Configuration



9. On the **INSTALLATION SUMMARY** page, click **Root Password** to set the root password for the operating system.
10. On the **INSTALLATION SUMMARY** page, click **Begin Installation** to begin installing the operating system.

2.3.3 Install with TUI

This chapter describes the complete procedure for deploying a Compute Node from an ISO, configuring the management network as a bond in the TUI, and installing the management node on that compute node.

Procedure

1. Select the boot option.

   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.

- GUI method
- VNC method
- Text mode method

2. Enter the INSTALLATION SUMMARY wizard.

   After the installation boot screen loads automatically, the **INSTALLATION SUMMARY** page appears. On this page, focus on the following three configuration items: **Installation Destination**, **Software Selection**, and **Root Password**.

3. Configure Installation Destination.

   a) Click **Installation Destination**.

   b) Under **Local Standard Disks**, select a disk to use as the system disk.

   c) Keep **Storage Configuration** set to the default **Automatic**.

   d) Click **Done** in the upper-left corner to return.

4. Configure Software Selection.

   a) Click **Software Selection**.

   b) Under **Base Environment**, select **Compute Node** for the compute node mode.

   c) Click **Done** in the upper-left corner to return.

5. Configure the Root Password.

   a) Click **Root Password**.

   b) Enter and confirm the root password.

      The password must meet the strength requirements indicated by the prompt.

   c) Click **Done** to return.

6. Start the installation and wait for it to finish.

   Verify that all items marked with red boxes are configured (no yellow warning bars remain), and click **Begin Installation** in the lower-right corner to start the installation.

   The installation completes automatically and requires no manual intervention. After the installati on finishes, the system reboots and the TUI appears.

7. Enter the TUI main screen and complete root authentication.

   The TUI is divided into four information areas:

   - **NODE INFO**: Node uptime, load, CPU, and memory usage.
   - **DISK INFO**: Disk mount points and capacity.

- **NIC INFO**: NIC names and IP addresses. The default NIC, such as ens3, is displayed at first boot.
- **VM INFO**: List of virtual machines on the current node.

The shortcut keys are shown at the bottom of the screen:

- **<F2> System Configuration**: Enter system configuration.
- **<F12> Reboot/Poweroff**: Reboot or shut down the system.

a) Press **F2** to enter system configuration. The **ROOT PASSWORD** authentication dialog box appears.

b) Enter the root password and press **Enter** to confirm.

8. Open the Configure Networks settings.

After authentication, the **System Configuration** menu appears.

a) Use the arrow keys to select **Configure Networks**, and then press **Enter**.

b) In the **CONFIGURE NETWORKS** dialog box that appears, select the physical NIC to configure (for example, ens3).

c) Switch the keyboard to **uppercase**.

d) Press the **A** key to trigger the **Link <A>ggregation** operation.

> **Note:**
> Make sure the keyboard is in **uppercase** mode. Otherwise, pressing **A** will not start the bond creation process.

9. Select the bond mode.

a) In the **CONFIGURE BOND** dialog box that appears, select a bond mode.

b) This chapter uses **Active-Backup** as an example. Press **Enter** to confirm.

10. Name the bond and select slave NICs.

The **CONFIGURE BOND SLAVE** dialog box appears. Configure the following settings.

a) In **BOND NAME**, enter a name for the bond, for example, bond0.

b) In **Select SLAVES**, use the arrow keys to navigate to the target NIC and press **Space** to select it.

A selected NIC is marked with `[X]`.

c) After selecting all required NICs, press **Enter** to confirm.

11. Confirm the bond configuration.

The **CONFIGURE BOND NETWORK CONFIRM** warning dialog box appears with the following message:

```
Be careful! Changing the network configuration may cause failure!
If you know what you are doing, then enter 'Y' to continue!
```

a) Switch the keyboard to **uppercase**.

b) Press the **Y** key to confirm the configuration.

c) When the operation is successful, the **SUCCESS** dialog box appears with the message `Network configuration is changed successfully!`. Press **Enter** to close it.

12. Select the bond to configure its IP.

You return to the **Configure Networks** dialog box.

a) In the NIC list, select the bond created in the previous step (for example, bond0).

b) Press **Enter** to enter the IP configuration.

13. Configure the management network IP address.

a) In the **CONFIGURE NETWORK** dialog box, enter the following fields as shown in the example:

- **VLAN ID**: Enter the VLAN ID of the management network. Leave it empty if no VLAN is used.
- **IP ADDR**: Enter the management network IP address, for example, *172.26.215.195*.
- **NETMASK**: Enter the subnet mask, for example, *255.255.0.0*.
- **GATEWAY(DEF)**: Enter the default gateway, for example, *172.26.0.1*.

b) Press **Enter** to confirm.

14. Confirm the IP configuration.

a) Press **Enter** to submit. The **CONFIGURE NETWORK CONFIRM** confirmation dialog box appears.

b) Switch the keyboard to **uppercase**.

c) Press the **Y** key to confirm.

d) The **SUCCESS** dialog box appears. Press **Enter** to close it.

15. Open the Install Nologo option.

a) Press **Esc** to return to the **System Configuration** menu.

b) Use the arrow keys to select **Install Nologo**.

c) Press **Enter**.

16. Confirm the management node installation.

The **INSTALL Nologo** warning dialog box appears with the following message:

```
This process will install a brand new Nologo, please confirm
```

a) Type the full string `install nologo`.

b) Press **Enter** to confirm.

17. Wait for the management node installation to complete.

The screen enters a waiting state and displays the message `Please wait...`. Wait for the installation to complete.

18. After the installation is complete, view the TUI main screen.

After the management node installation completes, the information on the TUI main screen changes as follows:

- **DISK INFO**: New mount points such as */var/lib* and */usr/local* are added.
- **NIC INFO**: A new `docker0` interface (related to the management service) is displayed, along with the configured `bond0` and its IP address (for example, *172.26.215.195*).

19. Verify the management node status from the command line.

a) Press **Ctrl + Alt + F2** to switch to the command line login screen.

b) Log in as the root user.

c) Run the following command to query the management node status:

```
[root@localhost ~]# cloud-ctl status
```

When the management node is running, the expected output is as follows:

```
MN status: Running
UI status: Running
```

2.4 Set Up Management Node HA

Management node HA improves the availability of the NexaSphere management plane. After management node HA is configured, if the active management node becomes unavailable, the standby management node can continue to provide management services.

Before setting up management node HA, prepare two servers on which the management node service has been installed, and plan the management network, business network, and VIP in advance. You can complete the setup from the user interface, or follow the CLI workflow to install the operating system, configure networks, and install the HA suite in sequence.

2.4.1 Set Up with the UI

Prerequisites

- You must install management node services on the selected two servers. For installation instructions, see *Install the Management Node*.
- Ensure both management nodes have gateway addresses on the same network segment and have consistent NIC names.
- Ensure both management nodes are running the same version.
- Ensure the operating system type and architecture are supported and both management nodes run compatible operating system types.
- Ensure both management nodes use the same license type.

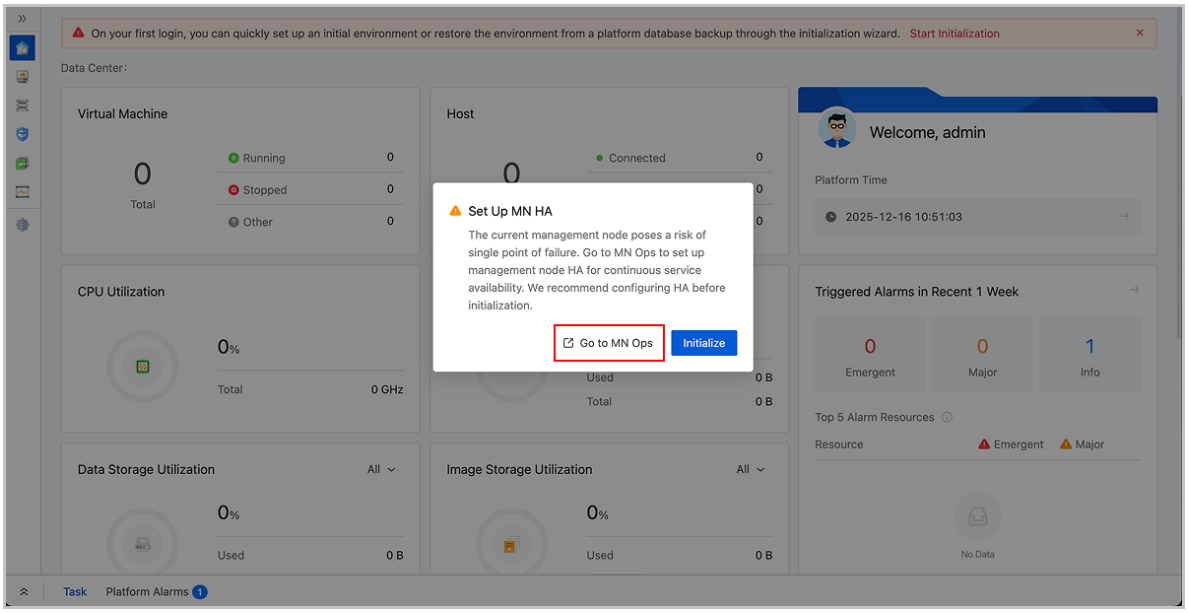
Procedure

1. Log in to the UI management interface of either management node.

   For a new environment that has not been initialized:

   1. In the **Welcome to Initialization Wizard** dialog, click **Next**. The system automatically detects the management node HA status.
   2. In the **Set Up MN HA** dialog, click **Go to MN Ops**.
   3. A new browser tab opens displaying the MN Ops page.

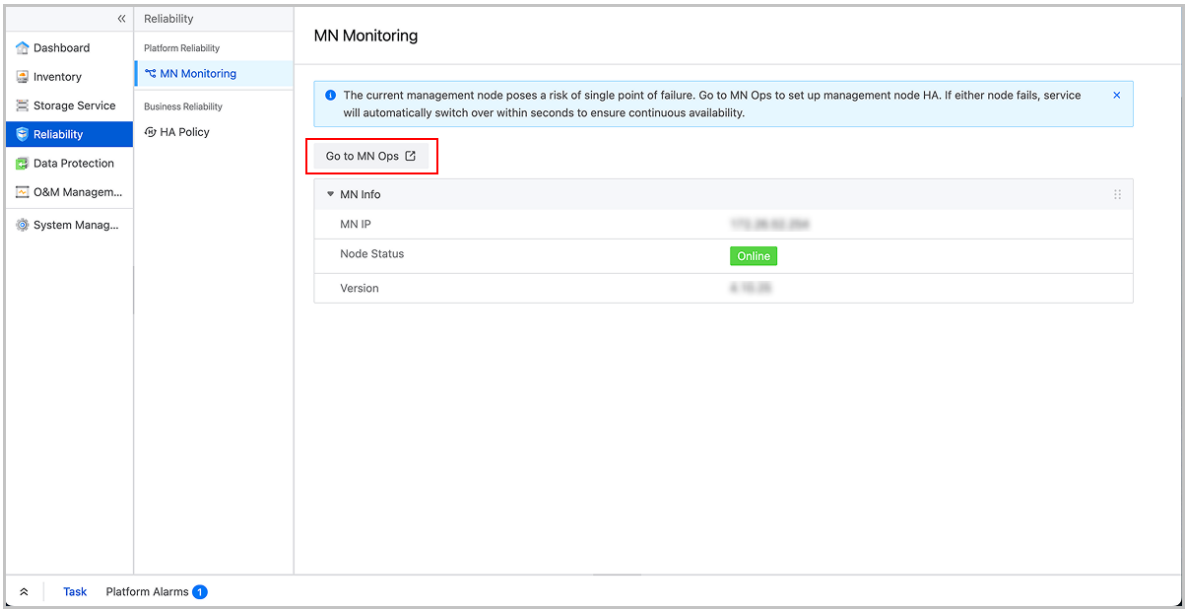
   Figure 2-21: Go to MN Ops from the Initialization Page

   For an existing environment:

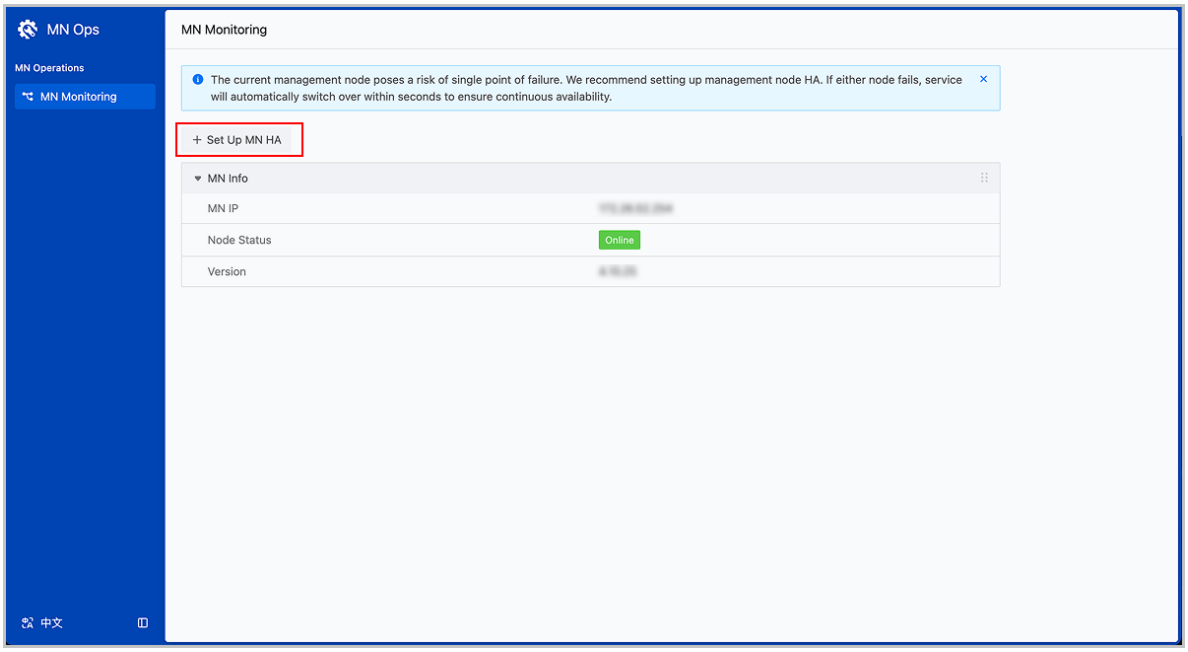
1. In the navigation pane, choose **Reliability** > **MN Monitoring**.
2. On the **MN Monitoring** page, click **Go to MN Ops**.
3. A new browser tab opens displaying the MN Ops page.

Figure 2-22: Go to MN Ops from the MN Monitoring Page



2. On the **MN Monitoring** page in MN Ops, click **Set Up MN HA**.

Figure 2-23: MN Monitoring



3. In the **Set Up MN HA** dialog, complete configuring MN, reviewing configuration, and setting up MN HA.

   a) For **Configure MN**, set the following parameters:

   - **VIP**: Configure the VIP for accessing the management interface of the management node environment.

     Note:
     Ensure the VIP is in the same network segment as the active management node IP.

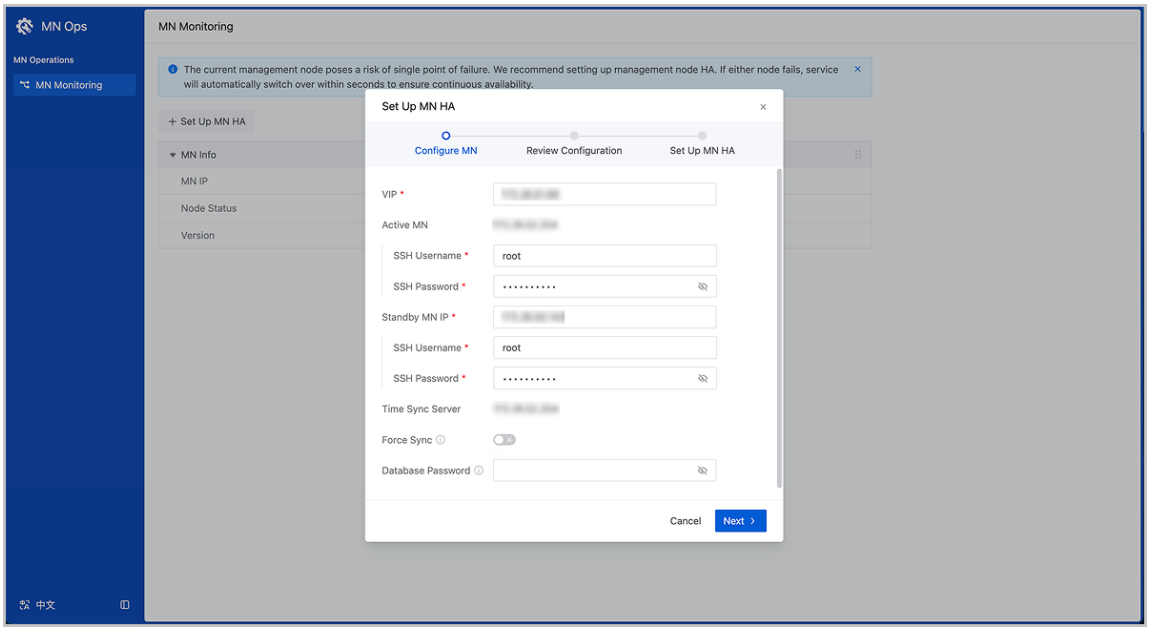
   - **Active MN IP**: Display the IP address of the active management node. Specify the SSH username and password for the active management node.
   - **Standby MN IP**: Specify the IP address, SSH username, and SSH password for the standby management node.

     Note:

     - Ensure the standby management node IP is in the same network segment as the active management node IP.
     - When setting up management node HA, the management node password can contain English letters, numbers, and these special characters: `-=[];,./~!@#$%^&*()_+|{}:<>?\``. If your password does not meet this rule, update it before setting up management node HA.

   - **Time Sync Server**: Uses the active management node IP as the time synchronization server by default.
   - **Force Sync**: Disabled by default. If the databases between active and standby management nodes cannot synchronize automatically, run the zsha2 installation command forcibly on the active management node.
   - **Database Password**: If left blank, the initial database password will be used by default. If the password has been changed, enter the new password here.

Figure 2-24: Configure MN



b) Click **Next**.

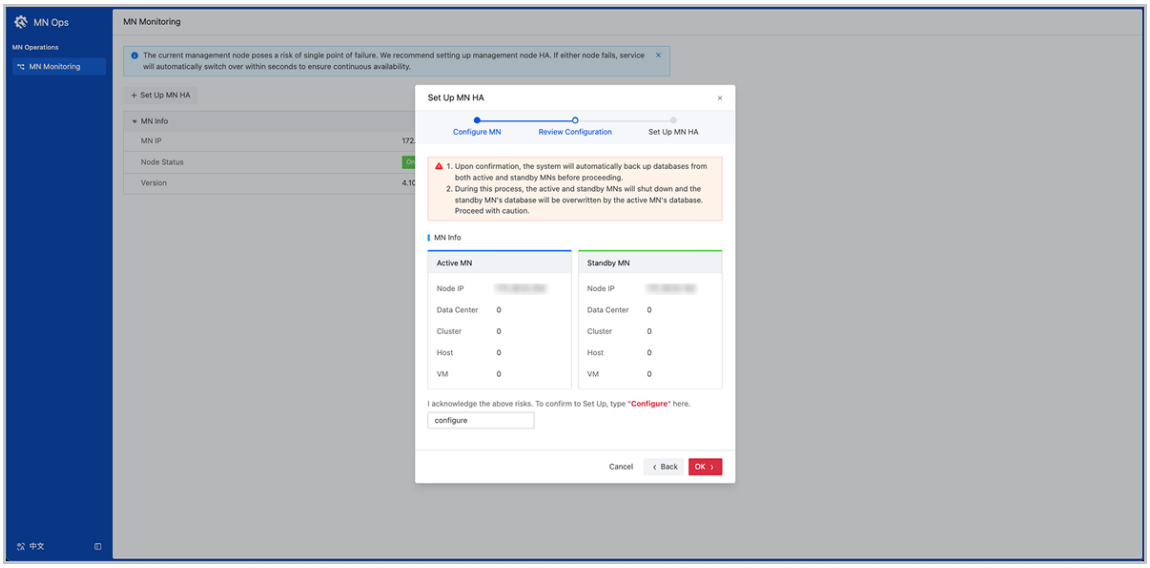
The system automatically checks if the requirements for addition are met.

c) In the **Review Configuration** step, check the resource information of both active and standby management nodes.

Note:

1. Upon confirmation, the system will automatically back up databases from active and standby management nodes before proceeding.
2. During this process, the active and standby MNs will shut down and the standby MN's database will be overwritten by the active MN's database. Proceed with caution.
3. If you find that the standby management node contains more resources during the resource comparison, and you want to preserve these resources, we recommend that you select the **Switch Active/Standby MNs** checkbox. When selected, this option switches the roles of active and standby management nodes, using the resource-rich standby management node as the active management node to set up management node HA.

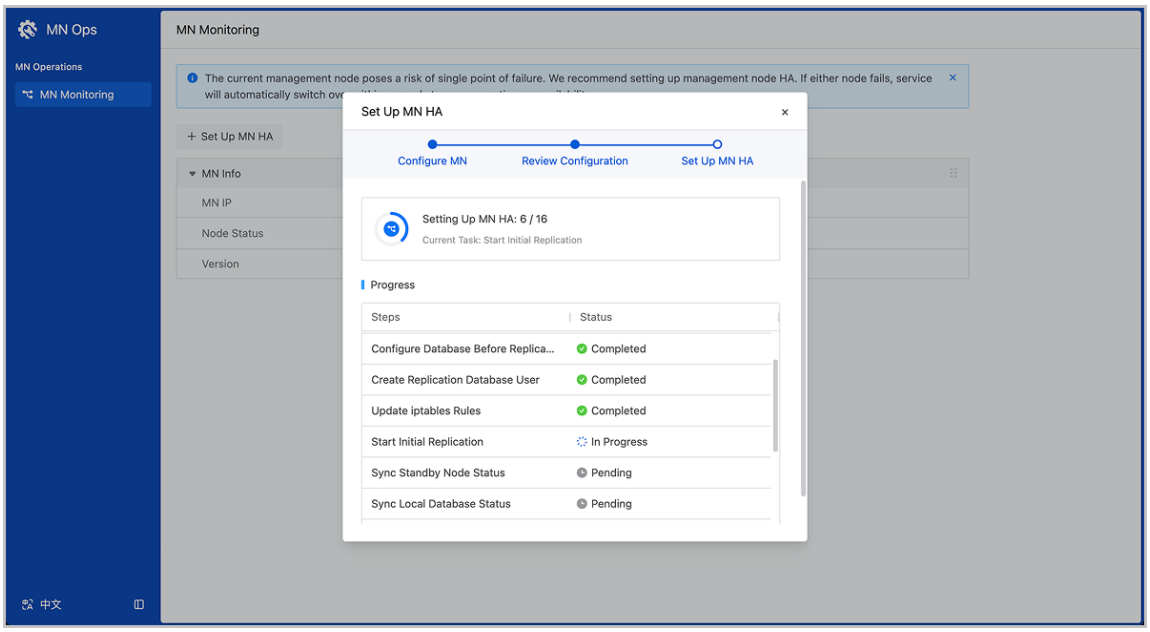
Figure 2-25: Review Configuration



d) Enter the confirmation information. Then, click **OK**.

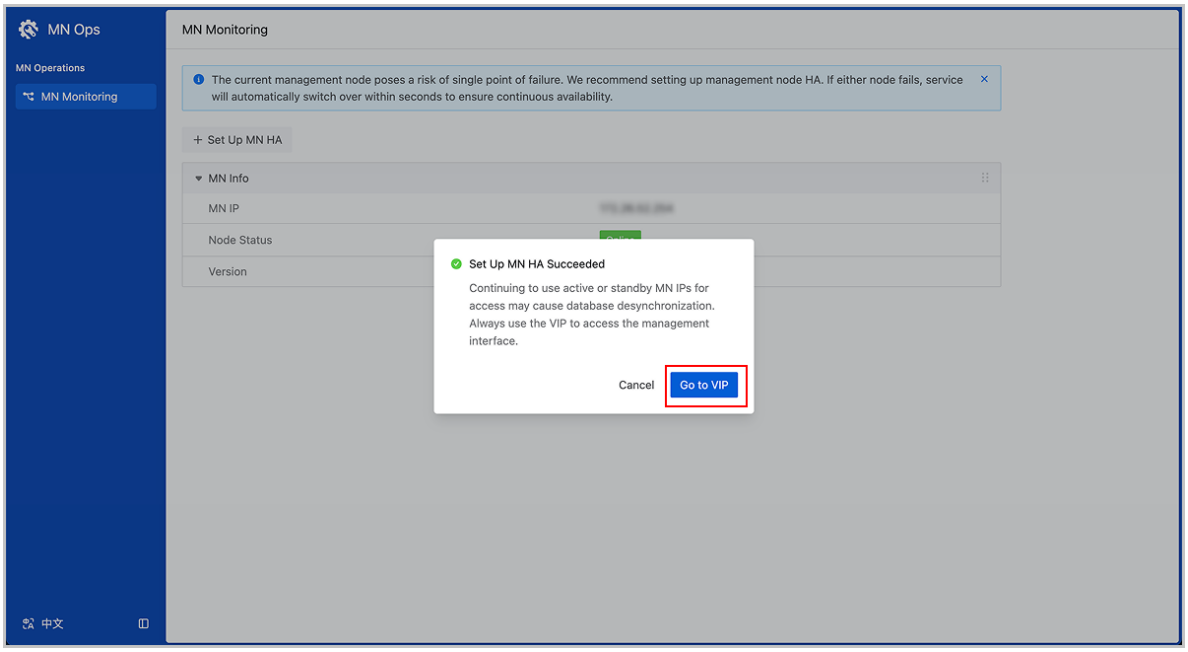
e) In the **Set Up MN HA** step, view the addition progress and task results.

Figure 2-26: Set Up MN HA



4. After the addition completes successfully, click **Go to VIP** to access the platform management interface through the VIP.

Figure 2-27: Go to VIP



2.5 Log in to NexaSphere

The platform uses HTTPS protocol by default and automatically redirects to port 443. You can simply enter the management node IP address in your browser to access the UI.

Prerequisites

- For a better experience, we recommend using Chrome 67 or later version with minimum screen resolution of 1280 × 900 px.
- If the web page fails to load, check whether both the management node and UI service are running properly.

Procedure

1. In a browser, enter the management node IP address in this format: *https://management_node_ip*.

   For a dual management node environment, use the VIP to access the UI.

2. Enter your username and password.

   On your first login, the default username is **admin** and the initial default password is **password**.

3. Click **Log In**.

What's next

- For more information about login policies, HTTPS/HTTP access, unified authentication, and certificates, see *Security*.
- After the first login, continue to initialize NexaSphere based on your deployment objective.

2.6 Initialize NexaSphere

After you log in to NexaSphere for the first time, you can use the initialization wizard to create basic resources so that the environment is ready for adding hosts, configuring storage, and creating virtual machines.

You can initialize the environment through the following methods:

- Initialize manually
- Restore from backup data

Notes during the initialization:

- If you exit the Wizard midway, **you will not be able to enter the Wizard again**.
- If no resources are created during the initialization, you can initialize the environment again from the Dashboard page.
- If you accidentally exit the Wizard, you can delete the data center and re-enter the initialization wizard from the Dashboard page.

> **Note:**
> Deleting a data center will delete clusters, hosts, networks, and data storage resources, along with all subordinate resources of each resource. Proceed with caution.

2.6.1 Initialize Manually

You can follow the wizard to initiate a new environment by creating necessary resources, including data center, cluster, host, data storage, image storage, image and distributed port group.

Procedure

1. In the **Welcome to Initialization Wizard** dialog, choose **Initialize Manually**.
2. Click **Next**.
3. On the **Initialize Manually** page, create a data center.
   a) Enter a name and description for the data center.
   b) Click **OK**.
4. Create a cluster.

a) Enter a name for the cluster.

b) Select a CPU architecture for the cluster. Options include x86_64 and aarch64.

c) Click **OK**.

5. Add a host.

a) Enter a name for the host.

b) Enter the host IP address, SSH port, SSH username, and SSH password.

> **Note:**
> The host architecture must match the cluster's CPU architecture, otherwise host addition to the cluster will fail.

c) Click **OK**.

The addition process may take several minutes to complete.

6. Add a data storage.

a) Enter a name for the data storage.

b) Select a data storage type. Options include local storage, nSDS distributed storage, and ZHPS distributed storage.

- Local storage: Set mount path and host disk configurations.

  > **Note:**
  > - Avoid system directories such as `/`, `/dev/`, `/proc/`, `/sys/`, `/usr/bin`, and `/bin` for mount path. Using system directories might cause the hosts unable to work properly.
  > - Host disk configuration will format the selected disks and completely erase all partitions, file systems, and data on the disk.

- nSDS distributed storage: Set whether to enable key authentication, add monitoring node, and specify image cache pool, storage pool, and storage network.

  > **Note:**
  > Make sure the key authentication configuration is consistent with the platform configuration. Otherwise, virtual machine creation may fail.

- ZHPS distributed storage: Enter IP address, port, username, and password. Add a storage pool.
- ZBS distributed storage: Add MDS nodes and Enter the pool name of the storage pool.

c) Click **OK**.

7. Add an image storage.

a) Enter a name for the image storage.

b) Select the image storage type. Options include standalone image storage and distributed image storage.

- Standalone image storage: Enter the image storage IP address, SSH port, username, and password. Complete the disk configurations, including addition method, mount path, whether to retrieve existing image, image sync network, and data network.



Note:

- If you select free disk for addition method, this configuration will format the selected disks and completely erase all partitions, file systems, and data on the disk.
- Avoid system directories such as `/`, `/dev/`, `/proc/`, `/sys/`, `/usr/bin`, and `/bin` for mount path. Using system directories might cause the image storage unable to work properly.

- Distributed image storage: Configure the monitoring node, including monitoring node IP address, SSH port, username, and password. Specify an image storage pool UUID as needed. If not specified, the platform creates one automatically.

c) Click **OK**.

8. Add an image.

a) Select the image path. Options include URL and local file.

b) Enter a name for the image.

c) Select the image type. Options include system image and disk image.

d) Select the image format. Options include qcow2, iso, vmdk, and raw.

e) Click **OK**.

9. Create a distributed port group.

a) Enter a name for the distributed port group.

b) Select the distributed switch for the distributed port group. Options include default switch and new switch.

- Default switch: Automatically created based on host configuration when adding the first host to a cluster. Primarily used for host management network. You can reuse this switch for both management network and business networks.

- New switch: To realize network separation between management network and business network, you can create a distributed switch as needed. This option is recommended for production environment. If you select this option, you need to configure aggregated interface configuration, uplink name, bond mode, hash policy, and host NC.

c) Select the VLAN type. Options include none and standard VLAN.

You need to specify a VLAN ID when selecting standard VLAN.

d) Choose whether to enable the DHCP service.

When enabled, you need to configure the network range method (IP range or CIDR), gateway, IP allocation policy, and DHCP IP.

e) Enter a DNS.

f) Click **OK**.

10. The **Manual Initialization Complete** dialog appears. Click **OK**.

You have completed the necessary resource creations for environment initialization. You may now start using the platform.

2.6.2 Restore from Backup Data

In the event of a disaster in the local data center, you can rely on the platform's database backup data to rebuild the data center and restore business. You can quickly restore the whole environment from an existing platform database backup using the initialization wizard.

Prerequisites

You have at least one valid platform backup.

Procedure

1. In the **Welcome to Initialization Wizard** dialog, choose **Restore from Backup Data**.
2. Click **Next**.
3. In the **Restore from Backup Data** dialog, set the following parameters to complete server configurations:

Parameter	Description
Backup Storage IP	Enter the IP address of the backup storage that stores the platform backup data.
URL	URL of the backup storage.
SSH Port	SSH port of the backup storage. Default is 22.

Parameter	Description
Username	Username of the backup storage.
Password	Password of the backup storage.

4. Click **Test Connection** to validate network connectivity to the backup storage.
5. After successful connection, set the following parameters to complete backup data configurations:

Parameter	Description
Database Backup	Select platform backup data.  **Note:** To avoid management node start failure due to insufficient licensed quota, select the appropriate data to restore. You can click **Update License** to renew your licensed quota.

6. Review the configuration and click **OK**.
7. In the **Restore Platform Database** dialog, enter the database root password and click **OK**.

 **Note:**
Restoring the platform database requires a management node restart, during which the management interface will be unavailable. This process takes a few minutes and does not affect your resources.

8. After the platform database is successfully restored, click **Log in Again**.

What's next

After platform database recovery, all resources on the platform will be recovered to the state at which the backup is created. Click **Scan Backup Data** in the local backup storage to obtain real-time backup data.

3 Licensing

3.1 Licensing Modes and License Rules

NexaSphere supports two licensing modes: Dedicated License and Shared License. Shared License is centrally managed in .

Licensing Modes

NexaSphere supports the following licensing modes:

- **Dedicated License**: NexaSphere uses its own license quota. A license can be provided by using a USB Key that has a license preloaded, or by using a Request Key to apply for and upload a License package.

Preset License and Upgrade Inheritance Rules

- A newly installed environment includes a preset Enterprise Trial license that is valid for three months and has no CPU limit.
- An upgraded environment retains the Primary License used before the upgrade. The authorizat ion remains unchanged whether the environment used an updated license or the original preset license before the upgrade.
- If the Primary License is updated to an Enterprise license after the upgrade, the license page displays the actual purchased authorization.

3.2 Use a Dedicated License

3.2.1 Use a USB Key License

3.2.1.1 Install a USB Key License

Administrators can install a license by using a USB Key that has the license preloaded.

Context

When you enter the **Licenses** page, the system reads USB Key information. The USB Key information is usually displayed after approximately 2 seconds.

Procedure

1. Obtain the USB Key hardware that has the license preloaded.
2. Insert the USB Key into a management node. In a dual-management-node environment, insert it into either management node.

3. In the NexaSphere navigation menu, choose **System Management** > **License Management** > **Licenses** to go to the **Licenses** page.

   Check the USB Key status and make sure the status is **Normal**.

4. Click **Synchronize** on the left of **Upload License**.
5. Check the license status and make sure the status is **Valid**.

3.2.1.2 Update a USB Key License

After authorization is expanded or added, administrators can use the USB Key ID to apply for an authorization file and update the license on the license management page.

Context

When you enter the **Licenses** page, the system reads USB Key information. The USB Key information is usually displayed after approximately 2 seconds.

Procedure

1. Insert the USB Key into a management node.
2. In the NexaSphere navigation menu, choose **System Management** > **License Management** > **Licenses** to go to the **Licenses** page.

   Check the USB Key status and make sure the status is **Normal**.

3. Click **Copy** USB Key ID, use this ID to apply for authorization, and obtain the authorization file.
4. Return to the **Licenses** page, click **Upload License**, and upload the authorization file.
5. Check the authorization information and make sure all information is correct and the license status is **Valid**.

3.2.2 Use a Request Key License

3.2.2.1 Install a Request Key License

When Request Key licensing is used, administrators can install a license by using the compressed License package.

Procedure

1. Obtain the compressed License package that contains the required authorization.
2. Upload the License file to the management node where the NexaSphere virtual IP is located. In a single-node deployment, upload it to that management node.
3. Check the license status and make sure the status is **Valid**.

3.2.2.2 Update a Request Key License

After authorization is expanded or added, administrators can download the Request Key, obtain a new compressed License package, and update the license on the license management page.

Procedure

1. In the NexaSphere navigation menu, choose **System Management** > **License Management** > **Licenses** to go to the **Licenses** page.
2. Click **Download Request Key**.
3. Use the Request Key to apply for and obtain the compressed License package that contains the required authorization.
4. Upload the License file to any NexaSphere management node. In a single-node deployment, upload it to that management node.
5. Check the license status and make sure the status is **Valid**.

3.3 View License Status and Licensing Records

Administrators can view the licensing mode, license status, license quota, and licensing history on the license management page.

View License Status

In the NexaSphere navigation menu, choose **System Management** > **License Management** > **Licenses** to go to the **Licenses** page. The page displays the current licensing mode and license status.

- Dedicated License: The page displays the licensed quota and used quota. If a USB Key is used , the page also displays the USB Key status and KEY ID.
- Shared License: The page displays the quota used by this site, total licensed quota, available quota, license server address, and connection status. The available quota is the total licensed quota minus the quota used by all sites, including this site.

When a Shared License is used, click **Sync** to obtain the latest license information. License uploads are not supported on this page. To update the license, contact the administrator to update the license package in the license pool.

View Licensing Records

After a license is installed or updated, click the **Licensing Records** tab on the **License** page to view detailed licensing history.

Field	Description
Upload Time	The time when the license was uploaded to the platform.
Licenses	The name of the license or licensed module.
Licensing Mode	The source of the license, including USB Key, Request Key, or Shared License.
Authorized Item	The authorized item of the license.
Authorized Quota	The authorized quota provided by the license.
State	The current status of the licensing record.
Issued on	The time when the license was issued.
Expires on	The time when the license expires.

3.4 License Status and Exceptions

During license usage, the page might display different USB Key statuses, license statuses, and exception reminders.

USB Key Status

Status	Description
Normal (Ready)	The USB Key status is displayed as normal when only one USB Key is inserted into the management node and is functioning properly.
Removed (Missing)	After the USB Key is removed, the USB Key status is displayed as removed. The platform checks whether the current License has expired. • If the License has not expired, a buffer License is generated. The buffer duration varies based on the removal scenario: ◦ For the first removal, the buffer License duration is 30 days. ◦ For a later removal, if the interval since the last removal is less than 60 days, the expiration time of the previous buffer License is inherited. If the interval is 60 days or longer, the buffer License is updated to 60 days.

Status	Description
Abnormal (Abnormal)	In a dual-management-node environment, if one management node has one USB Key inserted and the other management node has one or more USB Keys inserted, the USB Key status becomes abnormal. The abnormal USB Key status does not affect authorized usage.
Fault (Fault)	In a dual-management-node environment, if one management node has multiple USB Keys inserted while the other management node has no USB Key or multiple USB Keys inserted, a USB Key fault occurs. In a single-management-node environment, if multiple USB Keys are inserted into the management node, a USB Key fault occurs. **Note:** When the USB Key status is Fault and the buffer License has expired, authorization cannot be used.

License Status

Status	Description
Valid	When an active platform license is added, the corresponding position indicates Valid.
Expired	If the platform license has expired, the corresponding position indicates Expired and the functionality is unavailable.

Exception Reminders

Exception	Description
License to Expire	When the remaining validity of the license is less than 15 days, or the 7x24 after-sales service is within 30 days of expiration, a banner appears after login to indicate that the license is about to expire.
License Expired	When the platform license has expired, the platform automatically redirects to the **License Management** page after login to remind you that the license has expired.

Exception	Description
	Note: If existing services on the platform are still running, do not reconnect hosts, image storage, or data storage to avoid affecting service operation.
Insufficient License Quota	• Dedicated License: When you update a license and the actual resource usage exceeds the licensed quota, most platform features are restricted. The restrictions are as follows: ◦ You can log in to the platform and obtain, refresh, and view resource status. ◦ Operations that reduce quotas are supported, such as deleting hosts, deleting clusters, modifying host status, and deleting virtual machines. ◦ License management operations are supported, such as downloading the Request Key, uploading licenses, and deleting licenses. ◦ All other operations are restricted. **Note:** • License authorization includes CPU sockets, CPU cores, and hosts. If actual resource usage exceeds the licensed quota for any authorization method, platform usage is restricted. • After the management node is shut down, it cannot be restarted. You can only log in to the management node console and update the license in command-line mode. • Shared License: When the available quota in the license pool is insufficient, you cannot create resources, but existing resources continue to run. Contact the administrator to add quota to the license pool or adjust the sites that use Shared Licenses.

Exception	Description
License Server Connection Exception	If a site using a Shared License loses its connection to the license server, the platform temporarily uses cached license information. Check network connectivity or contact the administrator to restore the connection.

4 Data Center, Cluster, and Host Management

4.1 Overview

This document describes how to manage data centers, clusters, and hosts in a NexaSphere environment, and how to maintain host status and host hardware devices.

Data centers, clusters, and hosts are core objects for organizing and providing virtualization resources in NexaSphere. A data center groups resources, a cluster pools host compute capacity, and a host provides the physical compute resources required to run virtual machines.

Table 4-1: Basic Object Descriptions

Object	Description
Data Center	A data center is the top-level resource namespace in NexaSphere. It organizes resources such as clusters, hosts, storage, networks, and virtual machines.
Cluster	A cluster is a collection of hosts in a data center. It pools compute capacity and hosts virtual machines.
Host	A host is a physical compute node managed by NexaSphere. It provides CPU, memory, hardware devices, and network connectivity for virtual machines.

4.2 Data Center Management

4.2.1 Create a Data Center

Context

You create a data center from the root node. After the data center is created, you can create clusters in the data center and add hosts to those clusters.

Procedure

1. Navigate to the root node and click **New Data Center**.
2. Set the basic information of the data center.

Option	Description
Name	Enter the data center name.
Description	Enter a description of the data center as needed.

3. Click **OK**.

4.2.2 Edit a Data Center

Procedure

1. Open the target data center page.
2. Choose **Actions** > **Edit Name and Description**.
3. Modify the data center name or description.
4. Click **OK**.

4.2.3 Delete a Data Center

Procedure

1. Open the target data center page.
2. Choose **Actions** > **Delete**.

   Warning:
   Deleting a data center also deletes clusters, hosts, networks, data storage, and the resources contained in those objects, such as virtual machines. Proceed with caution.
3. Confirm the deletion.

4.3 Cluster Management

4.3.1 Create a Cluster

Context

You create a cluster in a data center. Configure cluster parameters based on the host architecture , network plan, and the VM placement, availability, and runtime settings you plan to use.

Procedure

1. Use one of the following methods to create a cluster.
   - In the left navigation pane, right-click the target data center and choose **New Cluster**.
   - In the left navigation pane, select the target data center, and then choose **Actions** > **New Cluster**.
   - On the **Cluster and Host** tab of the target data center, click **New Cluster**.
2. Configure basic information.

Option	Description
Name	Enter the cluster name.
Description	Enter a brief description of the cluster.
Data Center	Select the data center in which to create the cluster.
CPU Architecture	Select the x86 or aarch architecture. The architecture of hosts added to the cluster must be consistent with the cluster CPU architecture.
DRS	Monitors the CPU and memory loads of all hosts in the cluster and dynamically adjusts virtual machine workloads on hosts based on the scheduling policy. DRS is disabled by default. After it is enabled, manual scheduling and automatic scheduling are supported. For more information about configuring DRS, see *Configure DRS*.

3. Configure advanced settings as needed.

Cluster Network

Option	Description
VDI Network	Configures a dedicated VDI network CIDR for the cluster. If it is not set or does not take effect, the management network is used by default.
Migration Network	Configures a dedicated network CIDR for VM migration. If it is not set or does not take effect, the management network is used by default.

Cluster Overcommi

Option	Description
CPU Overcommit Ratio	Controls the number of virtual CPUs that can be allocated to virtual machines. Total allocatable virtual CPUs = total physical CPU cores in the cluster x CPU overcommit ratio.

Option	Description
Memory Overcommit Ratio	Controls the virtual memory capacity that can be allocated to virtual machines. Total allocatable virtual memory = (physical memory capacity - reserved capacity) x memory overcommit ratio.

Host Settings

Option	Description
Host CPU Model Check	Checks whether the CPU model of the source host is consistent with the CPU models of hosts in the current cluster during VM live migration or host addition. This option is disabled by default. If it is enabled, the corresponding operation is blocked when inconsistency is detected.
ignore_msrs Option	Enables the ignore_msrs option for the host KVM kernel module. This option is disabled by default.
Host Zero Copy	Enables Zero Copy on host CPUs. This option is disabled by default. After it is enabled, data copy operations between kernel mode and user mode are reduced, CPU usage time is reduced, and virtual NIC performance is improved.
Huge Pages	Enables huge pages on hosts in the cluster and allocates huge page space to virtual machines. Huge pages are disabled by default. Each huge page is 2 MB.
Host Reserved Memory	Sets the memory capacity reserved on the KVM host to ensure that the host system can run properly. After huge pages are enabled in the cluster, at least 4 GB memory must be reserved on each host.

VM Settings

Option	Description
VM HA	Enables the automatic restart mechanism after a virtual machine is powered off. This option is enabled by default. For more

Option	Description
	information about VM HA and resource scheduling, see *VM HA Policy* and *Cluster DRS*.
VM Cross-Cluster HA	Allows cross-cluster HA migration in shared storage scenarios, including NFS, SAN, and distributed storage, when VM HA starts a virtual machine on another host or when a host enters maintenance mode.
VM CPU Mode	Sets the virtual machine CPU model to match the host CPU model, so that the virtual machine can inherit part or all of the host CPU features for different service scenarios. The default value is not set (None) for the x86_64 architecture and passthrough for the aarch architecture. For more information about configuring VM CPU features, see *Configure VM CPU Features*.
VM Hyper-V	Enables Hyper-V emulation for virtual machines. This option is disabled by default and is mainly used in nested virtualization scenarios for Windows virtual machines.
Video Card Type on Boot	Sets the default graphics card type used when a virtual machine starts. The type determines whether basic, high-definition, or high-performance video capability is provided. The default value is vga for the x86 architecture and virtio for the aarch architecture. The aarch architecture supports only virtio.
vNIC Multi-Queue Upgrading	Enables a virtual machine NIC to use multiple queues to send and receive network packets, improving network PPS and bandwidth performance.

4. Click **OK**.

4.3.2 Modify a Cluster

Procedure

1. Open the target cluster page.
2. Select what you want to modify.
   - To modify the cluster name or description, choose **Actions** > **Edit Name and Description**.

- To modify cluster network, overcommitment, host settings, or virtual machine settings, choose **Actions** > **Modify Configuration**, or open the **Advanced Settings** tab.

3. After the configuration is complete, click **OK**.

4.3.3 Delete a Cluster

Procedure

1. Open the target cluster page.
2. Choose **Actions** > **Delete**.



Warning:

Deleting a cluster deletes all hosts in the cluster. If local storage is attached to the cluster, all virtual machines and snapshots on the hosts are also deleted. Proceed with caution.

3. Confirm the deletion.

4.4 Host Management

4.4.1 Add a Host

Context

After a host is added to a cluster, NexaSphere can centrally manage its compute resources, connection status, and hardware devices.

Procedure

1. Use one of the following methods to add a host.
   - In the left navigation pane, right-click the target cluster and choose **Add Host**.
   - In the left navigation pane, select the target cluster, and then choose **Actions** > **Add Host**.
   - On the **Host** tab of the target cluster, click **Add Host**.
2. Configure basic information.

Option	Description
Name	Enter the host name.
Description	Enter a brief description of the host.
Cluster	Select the cluster to add the host to.

Option	Description
Tag	Associate one or more tags as needed to identify different hosts. For more information about tags, see *Resource Tagging and Search*.

3. Configure host information.

Option	Description
Addition Method	Select whether to add a single host or multiple hosts. When adding multiple hosts, ensure that all hosts use the same SSH configuration.
IP Address	Enter the host IP address or IP address range based on the selected addition method.
SSH Port	Enter the SSH port of the host. The default value is 22.
SSH Username	Enter the host username. The default value is root.
SSH Password	Enter the SSH password for the specified username.

4. Configure other information as needed.

Option	Description
Scan Host IOMMU Setting	Enables or disables IOMMU (Input/Output Memory Management Unit). IOMMU is used for scenarios such as external device passthrough and virtualization. It is disabled by default on the x86 architecture. On the ARM architecture, IOMMU is called SMMU in the BIOS and is enabled by default.  **Note:** Before enabling IOMMU, ensure that the IOMMU option is enabled in the host BIOS.
Intel EPT Hardware Assist	Enables or disables Intel EPT Hardware Assist for Intel CPUs. This feature is enabled by default to improve CPU performance.

5. Click **OK**.

Result

After the first host is added to a cluster, NexaSphere automatically creates a default distributed switch, a default distributed port group, and a default Kernel adapter based on the host configuration to centrally manage the host management network. For more information about network resources, see *Network Model and Resources*.

4.4.2 Modify a Host

Procedure

1. Open the target host page.
2. Select what you want to modify.
   - To modify the host name or description, choose **Actions** > **Edit Name and Description**.
   - To modify the host IOMMU enablement status or Intel EPT Hardware Assist setting, choose **Actions** > **Modify Configuration**.
3. After the configuration is complete, click **OK**.

4.4.3 Access a Host

Context

You can access the host system through the Web terminal, SSH, or IPMI.

Procedure

Select one of the following methods to access the host.

- Web terminal access: Click the terminal window of the target host, or choose **Actions** > **Enter Web Terminal**.
- SSH access: Use an SSH client and the SSH information entered when the host was added. To modify SSH information, choose **Actions** > **Update SSH Information**.
- IPMI access: After IPMI is configured for the host, you can access the host by using IPMI management software. To modify IPMI information, choose **Actions** > **Modify IPMI Info**. IPMI management affects the host power status. For details, see *Host Status and Maintenance*.

4.4.4 Delete a Host

Procedure

1. Open the target host page.
2. Choose **Actions** > **Delete**.

Warning:

- If shared storage is attached to the cluster that contains the host, this operation also stops all virtual machines on the host. Virtual machines with HA enabled are automatically migrated to hosts with sufficient resources in the cluster and restarted.
- If local storage is attached to the cluster that contains the host, this operation also deletes all virtual machines and disks on the host. Proceed with caution.

3. Confirm the deletion.

4.5 Host Status and Maintenance

A host has three categories of status information: state, status, and power status. Each category indicates different host conditions and supported operations.

State

The state indicates whether the host is enabled, disabled, or in maintenance. It includes Enabled, Disabled, Pre-maintenance Mode, and Maintenance Mode.

- Enabled: You can create virtual machines on the host, start virtual machines on the host, or migrate virtual machines to the host.
- Disabled: You cannot create virtual machines on the host, start virtual machines on the host, or migrate virtual machines to the host.
- Pre-maintenance Mode: Before a host enters Maintenance Mode, it enters Pre-maintenanc e Mode first. Virtual machines on shared storage are automatically migrated to hosts with sufficient resources in the cluster. Virtual machines on local storage and virtual machines that cannot be migrated are forcibly powered off. After all virtual machines are migrated or stopped, the host enters Maintenance Mode.
- Maintenance Mode: Administrators can perform host maintenance, such as memory expansion , in this state. After the host enters Maintenance Mode, you cannot create virtual machines on the host, start virtual machines on the host, or migrate virtual machines to the host.

Operations that change the host state include enabling, disabling, entering Maintenance Mode, and exiting Maintenance Mode.

Status

The status reflects the network connection status between the host and the management node , including Connecting, Connected, and Disconnected. The host supports operations such as

creating or starting virtual machines only when the management node and the host are Connected .

When the network connection is unstable, the host status may change. Reconnecting a host changes its status.

Power Status

The power status reflects the power-on status of the host.

- IPMI Not Managed: IPMI information was not added when the host was added to the platform. The host power-on or power-off status cannot be obtained.
- Unknown: IPMI is configured for the host, but the power-on or power-off status cannot be obtained because IPMI is unavailable.
- Power On, Powering On, Power Off, and Powering Off: These power statuses can be obtained after IPMI is successfully configured for the host.

You can add or update IPMI information for the host by using **Update IPMI Info** only when the host power status is IPMI Not Managed or Unknown. Operations that change the host power status include adding or updating IPMI information, powering on, restarting, and powering off.

Powering off or restarting a host may power off virtual machines running on the host. To ensure service continuity and safety, you can put the host into Maintenance Mode in advance.

- If IPMI is configured for the host, soft shutdown is supported. You can stop services on the host before powering off the host.
- If IPMI is not configured for the host, the Shutdown operation powers off the host directly.

4.6 Host Hardware Device Management

NexaSphere supports viewing and managing hardware devices on hosts, including Host NUMA Topology, Block Devices, Host NICs, GPU Devices, USB Devices, and PCIe Devices.

Hardware device management helps administrators understand host hardware resource status and configure passthrough, virtualization partitioning, device attachment, and device maintenance based on service requirements.

4.6.1 Host NUMA Topology

Host NUMA topology: A pNUMA topology (physical NUMA topology) is the topology of the host NUMA nodes predefined by the CPU vendor based on the host NUMA architecture.

Definitions

- NUMA (Non-Uniform Memory Access): Non-uniform memory access (NUMA) is a computer memory design where the memory access time depends on the memory location relative to the CPU. Under NUMA, a processor can access its own local memory faster than non-local memory and thus improves VM performance.
- pNUMA Node: A pNUMA node (physical NUMA node) is a host NUMA node predefined based on the host NUMA architecture. It is used to manage the CPUs and memory of the host. A host can have one or more pNUMA nodes. A pNUMA node mainly consists of one or more physical CPU cores (pCPUs) and local memory.
- vNUMA Node: A vNUMA node (virtual NUMA node) is generated by passing-through associated pNUMA nodes via CPU pinning. It is used to manage the CPUs and memory of a virtual machine. A vNUMA node mainly consists of one or more virtual CPU cores (vCPUs) and local memory.
- vNUMA Topology: A vNUMA topology (virtual NUMA topology) is the topology of VM NUMA nodes generated by passing-through associated pNUMA nodes via CPU pinning.
- Local Memory: Local memory is the memory that a CPU (pCPU or vCPU) accesses through the Uncore iMC (Integrated Memory Controller) of the same NUMA (pNUMA or vNUMA) node. Compared with accessing non-local memory, accessing local memory has lower latencies.

How It Works

After a host is added to NexaSphere, you can view the host pNUMA topology and configure vNUMA for virtual machines running on the host based on the topology.

The virtual machine vNUMA configuration is implemented based on CPU pinning. CPU pinning strictly associates virtual machine vCPUs with host pCPUs and allocates specific pCPUs to the virtual machine. During vNUMA configuration, all vCPUs of the virtual machine are pinned to pCPUs, and each pCPU pinned to a single vCPU is in the same pNUMA node.

After vNUMA is configured, the associated host pNUMA node topology is exposed to the virtual machine, generating one or more vNUMA nodes that form the virtual machine vNUMA topology. Based on the vNUMA topology, the virtual machine vCPUs preferentially access local memory of their nodes.

pNUMA Topology

Open the **Overview** tab of the target host, and click **View pNUMA Topology** in the hardware overview information box. Host pNUMA topology information includes:

- All host pNUMA nodes and the virtual machines associated with each node.
- Total memory, which is the total local memory in the pNUMA node that the pCPU can directly access.
- Free memory, which is the free local memory in the pNUMA node that the pCPU can directly access.
- Total memory and free memory are both obtained from the real physical memory capacity of the pNUMA node.

You can also open the vNUMA topology of associated virtual machines.

4.6.2 LUN

Context

On the host **Hardware Device** tab, you can view the host's SCSI LUNs and NVMe LUNs.

Procedure

1. Navigate to **Inventory** > **VM and Host**.
2. Select the target host, and then click the host name to open the details page.
3. View SCSI LUNs.
    1. Go to **Hardware Device** > **LUN** > **SCSI Device**.
    2. View information about the LUN, including vendor, model, capacity, WWN, WWID, number of attached virtual machines, type, and source.
    3. Click the name of the target LUN, and then click **Paths** to view all available paths for the LUN and the status of each path.
4. View NVMe LUNs.

    Local NVMe PCIe disks on the host are displayed in the **Physical Disk** list.

    1. Go to **Hardware Device** > **LUN** > **NVMe Devices**.
    2. View information about the LUN, including vendor, model, WWN, capacity, WWID, and type.

4.6.3 Host NICs

On the host **Hardware Device** tab, select the **Physical NIC** tab to view and manage physical NICs and bonds on the host.

Physical NIC General Configuration

You can perform the following operations on host physical NICs:

Action	Description
Edit	Modifies the description of the physical NIC.
Modify IP Address	Modifies the IP address and netmask of the NIC as needed if the NIC has not been added to a bond or a distributed switch.

Physical NICs SR-IOV

Based on the SR-IOV specification, you can virtualize a physical NIC into multiple VF NICs and assign them directly to virtual machines. This provides I/O performance close to that of physical devices and reduces host CPU resource consumption.

Prerequisites

- Ensure that the physical NIC supports SR-IOV virtualization.
- Ensure that Intel VT-d/AMD IOMMU and SR-IOV are enabled in the BIOS of the host on which the physical NIC is installed.
- Ensure that the IOMMU status of the host on which the physical NIC is installed is **Available**.

Procedure

1. Navigate to **Host Details Page** > **Hardware Device** > **Physical NIC**.
2. Select the target physical NIC, and then click **Actions** > **Configure SR-IOV**.
3. In the **Configure SR-IOV** dialog, enable **SR-IOV Status** and specify the number of VF NICs to be partitioned.

Notes

- If the physical NIC is already configured in a bond, continuing to use SR-IOV may affect communication between VF NICs and vNICs. You are advised to partition the physical NIC through SR-IOV first and then configure the bond on the distributed switch. Select only one virtualized physical NIC on each host.
- If VF NICs are being used by virtual machines, disabling **SR-IOV Status** also detaches the related NICs from the virtual machines.
- Virtual machines that are powered on and have attached VF NICs do not support migration. Power off the virtual machine or detach the VF NICs before migration.

View and Maintain NIC LLDP Information

Link Layer Discovery Protocol (LLDP) identifies the physical switch port connected to a physical NIC of a specified distributed switch. Before you view peer device information, ensure that the NIC supports LLDP and that LLDP is enabled on the peer switch device.

Recommended Hardware Specifications

- Switches: Huawei switches, H3C switches, and Shengke switches.
- NICs: Intel 82599ES, Intel x710, Intel x722, and Mellanox CX4.

Supported OS Types

- x86: H84r and x86_KylinV10P3.
- ARM: arm_KylinV10P3 and H22e.

View peer device information:

1. Click the physical NIC name to open the details page.
2. Click **LLDP** to view peer device information. The following table describes the LLDP TLV units.

Modify the LLDP mode:

1. Click the physical NIC name to open the details page.
2. On the NIC details page, click **LLDP**.
3. On the **LLDP** tab, click **Modify**.
4. In the **Modify LLDP Mode** dialog, select the NIC LLDP mode from the drop-down list.
5. Click **OK**.

Mode	Description
Receive Only	This mode is selected by default. It parses and displays only peer LLDP information received by this port.
Send Only	Sends LLDP information only from this port and does not parse received LLDP information. In this mode, peer device information cannot be viewed.
Send and Receive	Parses and displays peer LLDP information received by this port, and sends LLDP information of this port to the peer device connected to it.

Mode	Description
Disabled	Does not parse received LLDP information or send local port LLDP information. In this mode, peer device information cannot be viewed.

Table 4-2: LLDP Field Definitions

Field	Description
Device ID	Chassis ID. The bridge MAC address of the sending device.
Port ID	The ID of the sending port.
Management Address	The management address of the sending port.
TTL	Time To Live, which indicates how long the local device information remains valid on neighboring devices.
Port Description	The description of the port.
System Name	The name that identifies the device.
System Description	The description that identifies the system.
System Capabilities	The main capabilities of the system and the capabilities that are enabled.
VLAN ID	The primary VLAN ID of the port.
Link Aggregation	Whether the port supports link aggregation and whether link aggregation is enabled.
MTU	Maximum Frame Size, which indicates the maximum frame length supported by the port and is derived from the configured maximum transmission unit (MTU) of the port.

View and Configure Bonds

If physical NICs were associated with a distributed switch for aggregation when the distributed switch was created, you can view the bond on the **Bond** tab of the corresponding host physical NIC.

Click the refresh button to view the latest information, including bond mode, bond status, rate, distributed switch association, IPv4 address, and creation time. For more information about creating and modifying bonds, see *Create a Distributed Switch* and *Modify Uplink Configuration*.

4.6.4 GPU Devices

On the host **Hardware Device** tab, select the **Physical GPU Device** or **vGPU Device** tab to view and manage physical GPU devices and vGPU devices on the host.

Physical GPU Devices

You can enable, disable, partition, or restore physical GPU devices as needed.

Operation	Description
Enable	Enables the physical GPU device so that it can be passed through directly to a virtual machine.
Disable	Disables the physical GPU device when you no longer want it to be passed through to other virtual machines.  **Note:** After the device is disabled, a physical GPU device that is being used by a virtual machine can continue to work properly until it is detached.
Virtualization Partitioning	Partitions a physical GPU device that is not passed through into vGPU devices of the specified specifications.
Restore Virtualization	Restores vGPU devices to a physical GPU device.

Before virtualization partitioning, ensure that the physical GPU device meets the following requirements:

- The physical GPU model supports virtualization partitioning.
- The physical GPU is not passed through to a virtual machine.
- Intel VT-d / AMD IOMMU is enabled in the host BIOS, and IOMMU support is enabled in the host kernel.
- The IOMMU readiness status of the host is **Available**.

The virtualization partitioning method varies slightly by physical GPU vendor.

- NVIDIA: Supports partitioning a single NVIDIA physical GPU based on the selected partitioning specification.
- AMD: Supports partitioning all AMD physical GPUs on the current host based on the selected partitioning quantity.

Before restoring virtualization, ensure that all vGPUs partitioned from the physical GPU are detached from virtual machines.

- NVIDIA: Before restoring NVIDIA vGPUs, ensure that all vGPUs created from this physical GPU are detached from virtual machines.
- AMD: Before restoring AMD vGPUs, ensure that all AMD vGPUs on the current host are detached from virtual machines.

vGPU Devices

You can enable or disable vGPU devices as needed.

Operation	Description
Enable	Enables the vGPU device so that it can be passed through directly to a virtual machine.
Disable	Disables the vGPU device when you no longer want it to be passed through to other virtual machines. After the device is disabled, a vGPU device that is being used by a virtual machine can continue to work properly until it is detached.

4.6.5 USB Devices

On the host **Hardware Device** tab, select the **USB Device** tab to view and manage USB devices on the host.

You can perform the following operations on USB devices as needed:

Operation	Description
Edit Device Name	Changes the USB device name to better meet service requirements.
Enable	Enables the USB device so that it can be passed through to a virtual machine.
Disable	Disables the USB device when you no longer want it to be passed through to other virtual machines. After the device is disabled, a USB device that is being used by a virtual machine can continue to work properly until it is detached.
Attach VM	Passes the USB device through directly to a virtual machine. Direct and Forward modes are supported.

Operation	Description
Detach VM	Detaches the USB device from a virtual machine. **Note:** This operation interrupts read and write services of the USB device. Proceed with caution.

Attach Modes

- Direct: Attaches a USB device from the host running the virtual machine to that virtual machine . Before migrating the virtual machine, detach the USB device first.
- Forward: Attaches a USB device from any host in the same data center as the virtual machine. You do not need to detach the USB device before migrating the virtual machine.

Attach Limits

- The same USB device can be passed through to only one virtual machine.
- A single virtual machine supports up to one USB 1.0 device, six USB 2.0 devices, and four USB 3.0 devices.
- A running virtual machine, or a stopped virtual machine that uses local storage, can use only available USB devices on the host running the virtual machine. Cross-host USB device attachment is not supported.
- A stopped virtual machine that uses shared storage can use multiple USB devices from a host in the same cluster as the virtual machine.

For more information about using USB devices on virtual machines, see *Create a Virtual Machine*.

4.6.6 PCIe Devices

On the host **Hardware Device** tab, select the **PCIe Device** tab to view and manage PCIe devices on the host.

You can enable passthrough for PCIe devices on the host. Passthrough devices help improve resource utilization and environment performance.

- **Passthrough PCI Devices**: Displays PCIe devices with passthrough enabled. These devices can be used by virtual machines.
- **All PCI Devices**: Displays all PCIe devices detected on the host. You can switch a device whose passthrough status is `Configurable` to a passthrough device.

Note:

- Before switching a PCIe device to passthrough, ensure that IOMMU is enabled on the host and that the IOMMU readiness status is Available.
- If a PCIe device is already attached to a virtual machine, it cannot be switched to passthrough. Detach it from the virtual machine and try again.

For more information about using PCIe passthrough devices on virtual machines, see *Create a Virtual Machine*.

5 Virtual Machine Management

5.1 Overview

NexaSphere VM management helps administrators prepare deployment resources, manage the virtual machine lifecycle, access and maintain virtual machines, configure virtual machine settings, perform VM migration, manage snapshots, and configure failure handling policies.

Before you create a virtual machine, prepare the required image storage, images, VM templates, or VM specifications. After the virtual machine is created, you can access it through the console or SSH, adjust its hardware and runtime configuration, install VMTools, migrate it to another host or data storage, and use snapshots to preserve short-term states.

For cross-platform migration scenarios, use Migration Service. For VM scheduling, VM high availability, backup and recovery, security groups, key providers, and trusted VM configuration, see corresponding manuals.

5.2 Prepare Images, Templates, and VM Specifications

Images, templates, and VM specifications preserve the software media, hardware configurat ion, and operating system initialization information required for VM deployment. Administrators can prepare image storage, add images, create templates, or create VM specifications before deployment to improve VM provisioning efficiency.

Image storage stores system images, disk images, and template files. VM templates are suitable for reusing virtual machine environments that have already been installed and configured. VM specifications are suitable for automatically configuring operating system parameters when you create virtual machines.

5.2.1 Image Storage Management

Image storage works as a repository for image templates that contain operating systems and for disk images. You can use image storage to share image files across data centers.

Before you create a virtual machine, prepare image storage and then add system images, disk images, or OVA templates to the image storage. The status of image storage affects operations such as image upload, synchronization, migration, and template creation.

5.2.1.1 Add a Standalone Image Storage

A standalone image storage stores image files as image slices and supports incremental storage.

Procedure

1. In the navigation pane, choose **Inventory** > **Image and Template**.
2. Right-click the target data center and choose **Add Image Storage**.
3. In the **Select Image Storage Type** dialog, choose **Standalone Image Storage**.
4. Click **Next**.
5. In the **Add Image Storage** dialog, configure the following options to complete basic configurations:

   a) Configure basic information.

Option	Description
Name	Image storage name.
Description	Enter a brief description for the image storage.
Type	Displays Standalone Image Storage.
Data Center	Current data center.

   b) Configure image storage information.

Option	Description
Image Storage IP	IP address of the image storage server.
SSH Port	Default port: 22. You can specify an SSH port.
Username	Default user: root. You can specify a regular user.
Password	Password of the user.

6. Click **Next**.

   The system tests connectivity to the image storage IP address.

7. Set the following parameters to complete storage configuration:

   a) Configure disk information.

Option	Description
Addition Method	Supports **Free Disk** and **Local Directory** addition methods.

Option	Description
Free Disk	When you select **Free Disk**, add free disks that are unmounted or unpartitioned on the image storage. **Note:** Configuring a free disk formats the selected disk and permanently removes all partitions, file systems, and data on the disk.
Mount Path	Absolute path of the storage mounted on the image storage. **Note:** Do not use `/`, `/dev/`, `/proc/`, `/sys/`, `/usr/bin`, `/bin`, or other system directories. Otherwise, image storage might work abnormally.
Retrieve Existing Images	Disabled by default. If enabled, the system retrieves existing image files from the mount path of this image storage.
Image Sync Network	Network CIDR used for image synchronization between standalone image repositories under the same management node. If left blank, the management network is used by default. If both the source and target image repositories are configured with this parameter, only the image sync network of the target image storage takes effect.
Data Network	Network CIDR used for data communication between compute nodes and image storage. If left blank, the management network is used by default.

b) Configure advanced settings.

Option	Description
Reserved Capacity for Image Storage	Default: 1 GB. Capacity reserved for image storage.
Image Blob Upload Concurrency	Default: 1. Concurrency for uploading image blobs when a single image is uploaded. Valid range: 1 to 16.

Option	Description
Image Blob Download Concurrency	Default: 1. Concurrency for downloading image blobs when a single image is downloaded. Valid range: 1 to 16.

8. Review the configuration and click **OK**.

5.2.1.2 Add a Distributed Image Storage

A distributed image storage stores image files on distributed block storage.

Procedure

1. In the navigation pane, choose **Inventory** > **Image and Template**.
2. Right-click the target data center and choose **Add Image Storage**.
3. In the **Select Image Storage Type** dialog, choose **Distributed Image Storage**.
4. Click **Next**.
5. In the **Add Image Storage** dialog, configure the following options:
   a) Configure basic information.

Option	Description
Name	Image storage name.
Description	Enter a brief description for the image storage.
Type	Displays Distributed Image Storage.
Data Center	Current data center.

   b) Configure image storage information.

Option	Description
Monitor Node	To add a monitor node, specify the management IP address, SSH port, username, and password of the monitor node.
Image Storage Pool	You can create a storage pool in the distributed storage cluster in advance and enter the UUID of the pool. If left blank, the system automatically creates an image storage pool.
Data Network	Network CIDR used for data communication between compute nodes and image storage. If left blank, the management network is used by default.

c) Configure advanced settings.

Option	Description
Reserved Capacity for Image Storage	Default: 1 GB. Capacity reserved for image storage.
Image Blob Upload Concurrency	Default: 1. Concurrency for uploading image blobs when a single image is uploaded. Valid range: 1 to 16.
Image Blob Download Concurrency	Default: 1. Concurrency for downloading image blobs when a single image is downloaded. Valid range: 1 to 16.

6. Review the configuration and click **OK**.

5.2.1.3 Manage Image Storage

You can modify image storage settings, view resource usage, clean up data, or delete image storage as needed.

Common Entries

You can manage image storage from the following entries:

- In the navigation pane, choose **Inventory** > **Image and Template**, and right-click the target image storage.
- Go to the target image storage page and click **Actions**.
- Go to the **Image Storage** tab of a data center and perform batch operations.

Common Operations

Operation	Description
Edit Name and Description	Modify the name and description of the image storage.
Modify Configuration	Modify the configuration of the image storage.
Update Password	Update the password of the image storage. After the password is updated, reconnect the image storage manually for the change to take effect.
View Usage	View image storage usage.

Operation	Description
Data Cleanup	Clean up permanently deleted invalid data and expired temporary data in the image storage. **Note:** After an image file is permanently deleted and the virtual machines created from the image are permanently deleted, data cleanup releases the storage space occupied by the image storage. Avoid write operations during data cleanup.
View Monitoring Data	View trends of image storage capacity, CPU, memory, disk resource usage, and NIC usage over time.
Delete Image Storage	Delete existing image storage. You can also delete image storage in bulk on the **Image Storage** tab of the data center.

5.2.1.4 Image Storage Status

Image storage has two status categories: state and status. Each category has a different meaning and supports different operations.

State

State indicates whether image storage is enabled or disabled by an administrator.

- Enabled: You can add images to this image storage and create virtual machines from the images on this image storage.
- Disabled: You cannot add new images. Existing images are not affected and can still be used to create virtual machines.

Status

Status indicates the network connection between image storage and the management node, including Connecting, Connected, and Disconnected. When the network is unstable, the image storage status might change. You can reconnect image storage to change its status.

5.2.2 Image Management

5.2.2.1 Add an Image

You can add system images or disk images in multiple formats to image storage. When you create a virtual machine for the first time, the system downloads the image to a data storage as an image

cache. You can add QCOW2, ISO, RAW, or VMDK image files to image storage by URL or local upload.

Procedure

1. In the navigation pane, choose **Inventory** > **Image and Template**.
2. Right-click the target image storage and choose **Add Image**.
3. In the **Add Image** dialog, configure the following options:

Option	Description
Image Storage	Current image storage.
Image Path	Add an image by URL or by uploading a local file.
URL	Supports HTTP, HTTPS, FTP, and SFTP URLs, or an absolute path on a standalone image storage.
Upload File	Select an image file that can be accessed by the current browser and matches the selected image format, and upload it directly. Interrupted local file uploads can be resumed.  **Note:** • The system automatically retries recoverable upload errors. If an upload is paused due to a network interruption and the page remains open, the upload resumes automatically after the network recovers. • If the browser retains permission to access the original local file after the page is refreshed or you sign in again, the system automatically resumes an upload that was paused due to an error. If the browser prompts you to allow access to previously selected files or folders, grant access as needed. • If the browser does not have permission to access the file or the permission has expired, the upload cannot resume automatically. Click **Pending Uploads** at the top of the page. In the **Pending Upload Tasks** panel, locate the task, click **Select to Continue**, and reselect the same local file.

Option	Description
	• A manually paused task does not resume automatically. In the task details, click **Continue**. • If the local file has been modified, moved, or deleted, the original upload task cannot resume. Start a new upload. • If an upload error occurs, the platform automatically retries the upload up to five times. If all five retry attempts fail, the upload task enters the Failed state and cannot be resumed. Start a new upload. • Click **View Upload History** to view local file upload tasks.
Name	Image name.
Description	Enter a brief description for the image.
Image Type	Image type, including system image and disk image.
Image Format	Available image formats vary by image type. Supported formats include *.qcow2*, *.iso*, *.raw*, and *.vmdk* files.

4. Review the configuration and click **OK**.

5.2.2.2 Synchronize Images

Context

Under the same management node, you can synchronize one or more images from a standalone image storage to a specified standalone image storage.

Image synchronization between image storage uses the management network by default. You can also configure a separate image synchronization network when you add an image storage or modify image storage configuration to reduce pressure on the management network. If both the source and target image storage are configured with image synchronization networks, only the image synchronization network of the target image storage takes effect during synchronization.

Procedure

1. In the navigation pane, choose **Inventory** > **Image and Template**.
2. Select the target image storage.
3. On the image storage details page, click **Image**.
4. On the **Image** tab, select the target image.

5. Click **Actions** > **Synchronize Image**.
6. In the **Synchronize Image** dialog, select the target standalone image storage.
7. Click **OK**.

5.2.2.3 Migrate an Image

Prerequisites

- Only migration between distributed image storage is supported.
- Before you migrate an image, make sure that the monitoring nodes of the two distributed image storage can communicate with each other.
- If an ISO image is migrated across data centers, virtual machines that use the ISO image must detach the image before they can start properly after the migration.

Procedure

1. In the navigation pane, choose **Inventory** > **Image and Template**.
2. Right-click the target image and choose **Change Image Storage**.
3. In the **Change Image Storage** dialog, select the target distributed image storage.
4. Click **OK**.

What's next

After migration, the original image data remains in image storage. You can manually clean up the data in the image storage. The cleanup cannot be undone. Proceed with caution.

5.2.2.4 Delete an Image

Prerequisites

Optional: The platform provides an image deletion protection mechanism. You can control deletion behavior by adjusting the deletion policy in system parameters. By default, the image deletion policy is delayed deletion, which keeps deleted resources for three days. Deleted resources are first moved to the recycle bin and then permanently deleted after the retention period expires.

Procedure

1. In the navigation pane, choose **Inventory** > **Image and Template**.
2. Select the target image.
3. On the image details page, click **Actions** > **Move to Recycle Bin**.

   To delete images in bulk, select target images on the Images tab of the image storage, and then click **Bulk Actions** > **Move to Recycle Bin**.

The delete button name changes according to the **Image Deletion Policy** setting. If the deletion policy is immediate deletion, the button name is Delete. If the deletion policy is delayed deletion or never deletion, the button name is Move to Recycle Bin.

4. After you acknowledge the risk information, click **OK**.

5.2.2.5 Export an Image

Select an image and export it. The exported image is displayed on the **Export List** tab of the root node. Exported images cannot be exported again.

Procedure

1. In the navigation pane, choose **Inventory** > **Image and Template** > **Image Files**.
2. Right-click the target image file and choose **Export Image**.
3. In the **Export Image** dialog, select an export method and click **OK**.



Note:

- For large image files, select Export Only. After export, you can view and download the file from the Export List page.
- For small and medium-sized image files, select Export and Download.

5.2.2.6 Export an OVA Template

Export a virtual machine in *.ova* format to a standalone image storage in the same data center. The exported OVA template does not include ISO CD-ROMs or GPU devices.

Prerequisites

- The virtual machine is shut down.
- VM encryption and disk encryption are disabled. An encrypted VM cannot be exported as an OVA template.

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. Right-click the target virtual machine and choose **Template** > **Export OVA Template**.
3. In the **Export OVA Template** dialog, select an export method and image storage, and click **OK**.

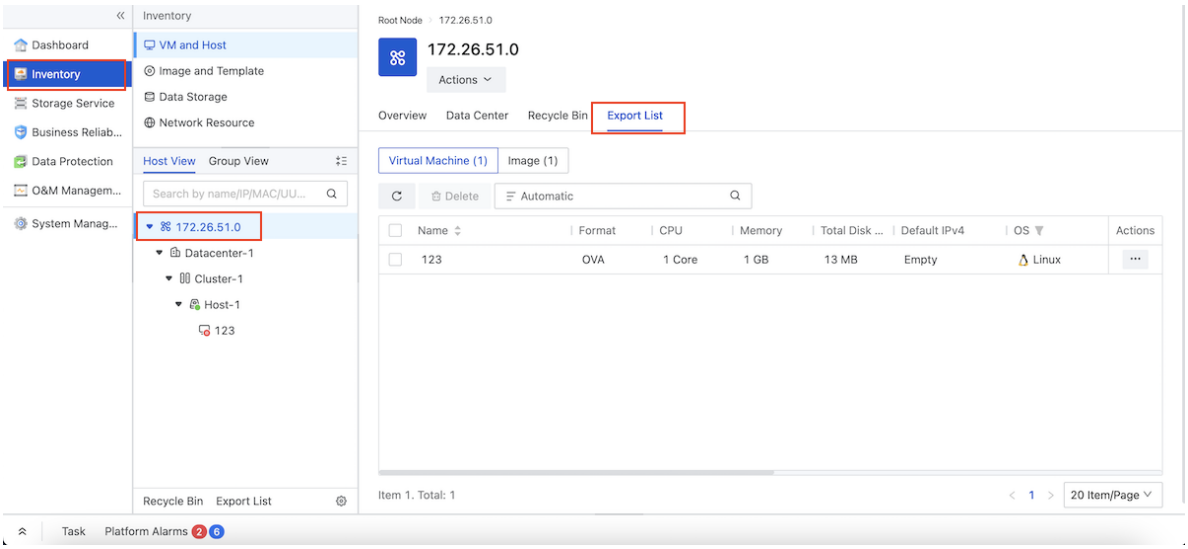


Note:

- For virtual machines with large files, select Export Only. After export, you can view and download the exported file from the Export List page.
- For virtual machines with small or medium-sized files, select Export and Download.

4. After export is complete, click **Inventory** > **Root Node**.
5. Click **Export List** to view the exported VM OVA template.

Figure 5-1: View Exported VM OVA Templates



5.2.3 VM Template Management

A VM template is a set of virtual machine operating system, application, and feature configurations. You can use VM templates to quickly deploy virtual machines with the same attributes in batches, improving efficiency and consistency.

5.2.3.1 Clone a VM to a Template

Clone a virtual machine to a template. The source virtual machine can continue to be used after it is cloned to a template.

Prerequisites

- The virtual machine power state must be Running, Paused, Stopped, or Crashed.
- The virtual machine must not have shared disks or RDM disks attached. If any are attached, detach them before cloning the virtual machine to a template.
- VM encryption and disk encryption must be disabled. An encrypted VM cannot be cloned to a template.

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. Select the target virtual machine and click **Actions** > **Clone** > **Clone to Template**.
3. In the **Clone Virtual Machine to Template** dialog, configure the following options:

Option	Description
Name	Enter a template name. **Note:** The name must be unique and cannot duplicate the names of existing virtual machines or templates.
Description	Enter a brief description for the template.
Data Storage	Location where the template is stored. By default, it is the same as the source virtual machine data storage.
Tags	Attach tags to identify different templates.

Note:
If the virtual machine has GPU, USB, or PCIe devices attached, except PCIe devices automatically attached by SR-IOV NICs, the cloned template does not include these peripheral devices. To keep these devices, you can convert the virtual machine to a template.

4. Review the configuration and click **OK**.

What's next

After you clone the virtual machine to a template, you can use the template to quickly deploy virtual machines with the same attributes.

You can modify the template without affecting normal use of the source virtual machine.

5.2.3.2 Convert a VM to a Template

Convert a shutdown virtual machine to a template. After the source virtual machine is converted to a template, it can be used only as a template and is removed from the virtual machine navigation tree. The template is displayed in the **Image and Template** > **Template File** navigation tree.

Prerequisites

- The virtual machine must be shut down.

- The virtual machine must not have shared disks or RDM disks attached. If any are attached, detach them before converting the virtual machine to a template.
- The virtual machine must not be associated with any backup plans. If it is associated with one, disassociate it before converting the virtual machine to a template.
- VM encryption and disk encryption must be disabled. An encrypted VM cannot be converted to a template.

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. Select the target virtual machine and click **Actions** > **Template** > **Convert to Template**.
3. Review the selected target and click **OK**.

What's next

After you convert the virtual machine to a template, you can use the template to quickly deploy virtual machines with the same attributes.

5.2.3.3 Convert a Template to a VM

Convert a template directly to a virtual machine. To edit template configuration or when you no longer need to keep the virtual machine as a deployment template, convert the template to a virtual machine.

Procedure

1. In the navigation pane, choose **Inventory** > **Image and Template**.
2. In the **Template File** list, select the target template and click **Actions** > **Convert to Virtual Machine**.
3. In the **Convert to Virtual Machine** dialog, configure the following options:

Option	Description
Name	Enter a virtual machine name.
Power Status	Choose whether to automatically power on the VM after conversion.
Location	Automatically assigned by default. Select the host or cluster where the virtual machine runs.

Option	Description
	• In local storage scenarios, you cannot specify the VM location . By default, the virtual machine runs on the host where local disks reside. • In non-local storage scenarios, you can specify the VM location .

4. Review the configuration and click **OK**.

What's next

After you convert the template to a virtual machine, you can adjust the virtual machine configuration as needed and then convert the virtual machine back to a template.

5.2.3.4 Create a VM from a Template

Create identical virtual machines from a template. You can customize hardware, software, and configuration for the new virtual machines based on the template.

Prerequisites

- The platform provides VM-related system parameters that globally control default platform behavior. Before creating virtual machines, you can customize VM-related system parameters to control default VM features.
- Before you create virtual machines from a template, make sure that a VM template already exists in the platform.
- To enable VM or disk encryption, make sure that an available default key provider is configured. If you use the Native Key Provider, back it up first. For more information, see *Encryption Keys and How Encryption Works*.

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. Right-click a valid parent resource of the virtual machine, such as a cluster or host, and choose **New Virtual Machine**.
3. In the **Select VM Creation Type** dialog, select **From Template**.
4. Click **Next**.
5. In the **Select Virtual Machine Template** dialog, select the target template and click **OK**.
6. In the **Create Virtual Machine from Template** dialog, configure the following options:

a) Configure template information.

Option	Description
Template	Select a template and create virtual machines from it.

b) Configure basic information.

Option	Description
Name	Virtual machine name.
Quantity	Number of virtual machines to create.  **Note:** When creating multiple virtual machines from a template, the following configurations will be cleared: • Added GPU, USB, and PCIe devices • Manually specified IP address
Group	Group to which the virtual machine belongs. If not specified, the default group is used.
Location	Host or cluster where the virtual machine runs.
OS	Virtual machine operating system. Mainstream operating systems are supported, including Linux and Windows.
HA	Mechanism for automatically restarting a virtual machine after an abnormal shutdown. Enabled by default.
VM Encryption	Encrypts runtime data written to disk, such as memory and GPU state, by using a key from the default key provider. Enabling this option automatically enables **Disk Encryption** for all disks. For more information, see *VM Encryption*.
Power Status	Whether to automatically power on the virtual machine after creation. Enabled by default.

c) Configure hardware information.

Option	Description
CPU	Supports adjusting the total number of cores and the number of cores per socket, and configuring CPU hot plug.
Memory	Supports adjusting memory size and configuring memory resource priority.
Disk	You can modify the following settings for an existing disk: • Adjust the disk cache mode and configure AIO acceleration. • **Disk Encryption**: Encrypts all data stored on the disk by using a key from the default key provider. You can encrypt an individual disk even when VM encryption is disabled. Disk encryption and shared-disk mode are mutually exclusive. Click **Add Hardware** > **Disk** to add a disk to the VM. You can customize the capacity and features of the new disk.
NIC	Supports adjusting NIC model, port group, multi-queue count, MAC address, IP address, and DNS assignment mode. You can click **Add Hardware** > **NIC** to add a new NIC to the virtual machine. You can customize the address and features of the new NIC.
CD/DVD Drive	Supports attaching an ISO image file to the virtual machine so that the virtual machine can boot from the ISO CD-ROM.
GPU Device 1	Attach a GPU device to the virtual machine. Physical GPU devices and vGPU devices are supported. Click **Add Hardware** > **GPU Device** to add a GPU device.
USB Device 1	Attach a USB device to the virtual machine. Direct connection and forwarding are supported. Click **Add Hardware** > **USB Device** to add a USB device. A single virtual machine supports at most one USB device.
PCIe Device 1	Attach a PCIe device to the virtual machine. PCIe devices switched to passthrough mode on the host are supported. Click **Add Hardware** > **PCIe Device** to add a PCIe device.

Option	Description
Other Hardware	The video card, sound card, and motherboard configuration in a template file cannot be modified.

d) Configure advanced settings.

General Options

Option	Description
Description	Displays the description recorded in the template by default. You can modify the virtual machine description.
Tag	Displays the VM tags recorded in the template by default. You can modify tags.
OS Attribute	Configures the VM operating system attributes. • Do Not Customize: Inherits the hostname, administrator password, workgroup or domain configurations from the template. • Apply a Specification: Select an existing VM specification to apply the system configuration specified in the specification. • Manually Customize: Customize a new VM specification.

7. Review the configuration and click **OK**.

What's next

Some virtual machine configurations depend on VMTools. After the virtual machine is created, install VMTools to make these configurations take effect. For more information about VMTools, see *Install and Use VMTools*.

5.2.3.5 Manage VM Templates

Modify VM templates as needed based on service scenarios, including editing the name and description, modifying configuration, and changing the owner. If you no longer need a template, delete it.

Common Entries

You can manage VM templates from the following entries:

- Go to the target template page and click **Actions**.

- Go to the **Image and Template** tab of the data center that contains the templates and perform batch operations.

Common Operations

Operation	Description
Edit Name and Description	Modify the name and description of a template. **Note:** The new name must be unique and cannot duplicate the names of existing virtual machines or templates.
Modify Configuration	Modify template basic information, total CPU cores, CPU cores per socket, memory size, or disk capacity. To make extensive changes to template configuration, convert the template to a virtual machine, adjust the configuration as needed, and then convert the virtual machine back to a template for later use.
Change Owner	Specify the owner of a template.
Delete VM Templates	Delete a single template, or delete templates in bulk on the **Image and Template** tab of the data center that contains the templates. **Note:** After a template is deleted, its template files, including NIC files, disk files, and database records, cannot be restored. Proceed with caution.

5.2.4 VM Specification Management

You can create and manage customization specifications for Windows or Linux virtual machine operating systems. A VM specification is a file that contains virtual machine operating system configuration.

VM customization specifications help avoid conflicts during batch deployment from templates, such as duplicated hostnames or SIDs.

5.2.4.1 Create a Windows VM Specification

Save specific Windows VM operating system settings in a customization specification, such as how the virtual machine participate in the network and hostname configuration.

Prerequisites

The following operating system versions support customization specification for virtual machines:

- Windows Server 2012
- Windows Server 2016
- Windows Server 2019
- Windows Server 2022
- Windows Server 2025
- Windows 10 Pro

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. Select the root node.
3. On the root node details page, click **VM Specifications**.
4. On the **VM Specifications** tab, click **New VM Specification**.
5. In the **New VM Specification** dialog, configure the following options:

Option	Description
Name	Enter a name for the VM customization specification.
Description	Enter a brief description for the VM specification.
Target VM OS	Select **Windows**.
Workgroup or Domain	Select how the virtual machine participates in the network. Options include workgroup and Windows domain server. For workgroup, you need to specify the workgroup name and choose whether to generate a new SID.  **Note:** The workgroup name must 1 to 15 characters in length and can contain letters, numbers, underscores (_), dots (.), and hyphens (-). The workgroup name cannot be all numbers.

Option	Description
	For Windows domain server, you need to configure the domain name, domain username, domain password, OU, and choose whether to generate a new SID. **Note:** • The domain name must be 1 to 26 characters in length and can contain letters, numbers, underscores (_), dots (.), and hyphens (-). • The domain username must 1 to 15 characters in length and can contain letters, numbers, underscores (_), dots (.), and hyphens (-). The domain username cannot be all numbers. • OU example: OU=MyOU,DC=MyDom,DC=MyCompany,DC=com
Specify Hostname	Select how to configure a VM hostname. Options include using the VM name and entering a name. **Note:** • When you use the virtual machine name as the hostname , the Windows VM name will be truncated if it exceeds 15 characters. • The following rules apply for specifying a Windows hostname: ◦ Length: 2 to 15 characters. ◦ Allowed characters: Uppercase and lowercase letters, numbers, and hyphens (-). ◦ No consecutive hyphens. You cannot start or end with a hyphen. The hostname cannot be all numbers. • When creating multiple VMs at once, a suffix (-1, -2, -3, etc.) is automatically appended to ensure uniqueness.
Administrator Password	Set the administrator password.

Option	Description
	Note: The password must contain 8 to 16 characters and at least two of the following character types: lowercase letters, uppercase letters, digits, and special characters.
Confirm Password	Enter the password again for confirmation.

6. Review the configuration and click **OK**.

5.2.4.2 Create a Linux VM Specification

Save specific Linux VM operating system settings in a customization specification, such as the hostname configuration.

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. Select the root node.
3. On the root node details page, click **VM Specifications**.
4. On the **VM Specifications** tab, click **New VM Specification**.
5. In the **New VM Specification** dialog, configure the following options:

Option	Description
Name	Enter a name for the VM customization specification.
Description	Enter a brief description for the VM specification.
Target VM OS	Select **Linux**.
Specify Hostname	Select how to configure a VM hostname. Options include using the VM name and entering a name. **Note:** • When you use the virtual machine name as the hostname , the Linux VM name will be truncated if it exceeds 60 characters. • The following rules apply for specifying a Linux hostname: ◦ Length: 2 to 60 characters.

Option	Description
	◦ Allowed characters: Uppercase and lowercase letters, numbers, and hyphens (-). ◦ No consecutive hyphens. You cannot start or end with a hyphen. • When creating multiple VMs at once, a suffix (-1, -2, -3, etc.) is automatically appended to ensure uniqueness.
Administrator Password	Set the administrator password. **Note:** The password must contain 8 to 16 characters and at least two of the following character types: lowercase letters, uppercase letters, digits, and special characters.
Confirm Password	Enter the password again for confirmation.

6. Review the configuration and click **OK**.

5.2.4.3 Manage VM Specifications

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. Select the root node.
3. On the root node details page, click **VM Specifications**.
4. On the **VM Specifications** tab, select the target VM specification.

Action	Description
Actions > **Edit Name and Description**	Modify the name and description of the VM specification.
Actions > **Modify Configuration**	Modify VM specification configuration, such as network participation mode, host name configuration method, and administrator password reset.
Actions > **Delete**	Delete the VM specification.

5.3 Deploy VM and Manage the VM Lifecycle

VM lifecycle management covers creating, importing, registering, cloning, and deleting virtual machines. Different deployment methods apply to VM files from different sources and to different operational scenarios.

Create a VM when you want to start from an image, a template, or a custom configuration. Import a VM to bring in an existing VM through an OVF file. Register a VM to bring existing VM configurat ion and data files under platform management. Clone a VM to quickly copy an existing VM environment.

5.3.1 Create a Virtual Machine

A new virtual machine that you'll be able to customize CPU, memory, storage, and network settings for the workload.

Prerequisites

- The platform provides VM-related system parameters that control default platform behavior. Before creating a VM, you can customize VM-related system parameters to control default VM characteristics.
- The platform provides default capacity and feature settings for VMs, such as **CPU**, **Memory** , **Disk 1**, **NIC 1**, **Other Hardware**, and **Advanced Settings**. You can quickly create a VM by using the default platform settings.
- You can create a VM without an OS image. You can add the disk 1 by creating a new disk. A CD/DVD drive is optional.
- You can create a VM without a NIC and configure NICs and their features later.
- To enable VM or disk encryption, make sure that an available default key provider is configured. If you use the Native Key Provider, back it up first. For more information, see *Encryption Keys and How Encryption Works*.

Procedure

1. In the navigation pane, right-click the target cluster, host, or image, and then click **New Virtual Machine**.
2. In the **Select VM Creation Type** dialog, select **New VM**, and then click **Next**.
3. In the **New Virtual Machine** dialog, configure the basic information as needed:

Option	Description
Name	VM name.
Quantity	Number of VMs to create.
Group	VM group. If not specified, the default group is used.
Location	Host or cluster where the VM runs.
OS	VM operating system. Mainstream operating systems are supported, including Linux and Windows.
HA	Automatic restart mechanism after an abnormal VM shutdown.
VM Encryption	Encrypts runtime data written to disk, such as memory and GPU state, by using a key from the default key provider. Enabling this option automatically enables **Disk Encryption** for all disks. Enable encryption during VM creation when possible to avoid the additional time required to convert existing data later. For more information, see *VM Encryption*.
Power Status	Whether to automatically power on the VM after it is created. By default, the VM is powered on automatically.

4. Configure the VM hardware settings as needed:

   a) CPU: Customize the VM CPU cores and features.

Option	Description
Cores	Number of CPU cores in the VM.
Cores per Socket	Number of CPU cores allocated per socket in the VM.
CPU Mode	Whether to match the VM CPU model with the host CPU model so that the VM can inherit some or all host CPU features for different workloads. The default value is None. For more information, see *Configure VM CPU Features*.
CPU Resource Priority	For critical VMs, set this option to High to give the VM higher CPU contention capability than VMs with Normal priority when the host is overloaded. For more information, see *Configure VM Resource Contention*.

Option	Description
CPU Clock Speed Limit	Controls the upper limit of host CPU resources that the VM can use. The valid range is 1% to 100%. Setting the value to 100% or leaving it blank means no limit.
CPU NUMA Binding	Binds VM CPUs (vCPUs) to host CPUs (pCPUs) and assigns specific pCPUs to the VM to improve VM performance. For more information, see *Configure VM Resource Contention*.
CPU Hot Plug	Whether the VM supports online CPU modification. This option is enabled by default. CPU hot plug and memory hot plug must be enabled or disabled together.  **Note:** Only some operating systems support hot plug. Click **View** to view the list of operating systems that support hot plug.
CPU Hypervisor Tag	Whether to enable the virtualization (hypervisor) flag for the VM. If disabled, applications skip VM virtualization environment detection. This option is enabled by default.

b) Memory: Customize the VM memory capacity and features.

Option	Description
Memory	VM memory capacity.
Memory Resource Priority	For critical VMs, set this option to High to give the VM higher memory contention capability than VMs with Normal priority when the host is overloaded. For more information, see *Configure VM Resource Contention*.
Memory Hot Plug	Whether the VM supports online memory modification. This option is enabled by default. CPU hot plug and memory hot plug must be enabled or disabled together.  **Note:** Only some operating systems support hot plug. Click **View** to view the list of operating systems that support hot plug.

c) Disk 1: Customize the storage location, capacity, and features of the VM system disk. You can create a new disk or select an existing system image.

Option	Description
Disk Encryption	Encrypts all data stored on the disk by using a key from the default key provider. You can encrypt an individual disk even when VM encryption is disabled. Disk encryption and shared-disk mode are mutually exclusive.
Storage Location	Data storage location of the VM disk. If not specified, the system allocates a data storage automatically.
Capacity	System disk capacity when creating a new disk.
System Image	When using a system image, select an image that contains the VM operating system. raw and qcow2 formats are supported.
Bus Type	Bus type of the VM disk device. Options include Virtio, IDE, Virtio SCSI, and SCSI. By default, Disk 1 uses Virtio on Linux and IDE on Windows and Other operating systems.
Provision Method	When SAN storage is selected as the storage location, select how disk storage space is allocated. The default value is Thin Provisioning. • Thin Provisioning: Allocates storage space based on actual usage for higher storage utilization. • Thick Provisioning: Pre-allocates the required storage space when the disk is created to provide sufficient storage capacity and ensure storage performance.
Caching Mode	Whether to use the host page cache for VM write operations and, if used, whether to return success only after data is written to the storage device. The default value is none, which means the host page cache is not used.
AIO Acceleration	Whether to enable asynchronous I/O (AIO) acceleration in the VM kernel. This option is disabled by default. **Note:**

Option	Description
	AIO acceleration can be enabled only when the caching mode is none.
QoS	Whether to set read/write bandwidth and IOPS limits for the disk. By default, no limit is set. For more information, see *Configure VM QoS*.

Disks other than Disk 1 are data disks, such as Disk 2. Click **Add Hardware** > **Disk** to add a data disk. A single VM supports up to 24 disks, including Disk 1.

- Supported data disk addition methods include creating a new disk, using a disk image, selecting an existing disk, and adding an RDM disk.
- When you create a disk or add one from a disk image, **Disk Encryption** follows the VM encryption state by default and can be changed as needed.
- When you select an existing disk, it retains its encryption state, which cannot be changed during attachment. The VM encryption state does not change the encryption state of the existing disk.
- RDM disks do not support disk encryption and are not automatically encrypted when VM encryption is enabled.
- By default, data disks use Virtio on Linux and Windows. For Other operating systems, data disks support only IDE.
- When the storage location is ZHPS distributed storage, only Virtio is supported as the bus type.
- When the disk storage location is nSDS distributed storage and the bus type is Virtio SCSI, the data disk can be set as a shared disk. This is not set by default.
- When the disk storage location is SAN storage, the bus type is Virtio SCSI, and the provision method is Thick Provisioning, the data disk can be set as a shared disk. This is not set by default.
- A shared disk cannot be encrypted. If **Disk Encryption** is enabled, disable it before configuring the disk as shared.
- When a disk is a shared disk, you cannot modify its bus type or QoS limit.

d) NIC 1: Customize the VM NIC address and features.

Option	Description
State	Whether to automatically enable the NIC when the VM powers on.

Option	Description
NIC Model	Set the NIC model. e1000, rtl8139, virtio, and SR-IOV are supported.
Port Group	Port group that contains the distributed switch port connected to the VM NIC.
NIC Queue Number	Uses multiple queues for the VM NIC to send and receive network packets, improving network PPS and bandwidth performance.
MAC Address	Specify a VM MAC address. By default, the system assigns a MAC address automatically.
Specify IP Address	Specify a VM IP address. By default, the system assigns an IP address automatically. If DHCP is disabled for the distributed port group where the VM NIC resides, you can use VMTools to specify an IPv4 address and automatically apply it to the VM. For more information, see *Install and Use VMTools*.
Assign DNS	Specify how DNS addresses are assigned. By default, the system assigns DNS addresses automatically.
Security Group	Associate security groups with the VM NIC to control east-west traffic. The smaller the value on the left of an associated security group, the higher its priority. **Note:** Configure security groups carefully to avoid rule conflicts between them.
QoS	Whether to set bandwidth limits for packets sent and received by the VM NIC. By default, no limit is set. For more information, see *Configure VM QoS*.

Click **Add Hardware** > **NIC** to add multiple NICs.

e) CD/DVD Drive 1: Mount an ISO system image file to boot the VM from a CD/DVD drive.

Click **Add Hardware** > **CD/DVD Drive** to add a CD/DVD drive. A single VM supports up to 3 CD/DVD drives.

f) GPU Device 1: Attach a GPU device to the VM.

Physical GPU devices and vGPU devices are supported. Click **Add Hardware** > **GPU Device** to add a GPU device.

g) USB Device 1: Attach a USB device to the VM.

Direct connection and forwarding are supported. Click **Add Hardware** > **USB Device** to add a USB device. A single VM supports up to 1 USB device.

h) PCIe Device 1: Attach a PCIe device to the VM.

PCIe devices on the host that have been switched to passthrough are supported. Click **Add Hardware** > **PCIe Device** to add a PCIe device.

i) TPM: Attach a virtual Trusted Platform Module (vTPM) to provide enhanced security for the VM operating system.

Before adding a TPM device, add an available key provider and set the VM boot mode to UEFI. Click **Add Hardware** > **TPM** to add a TPM device. A single VM supports up to 1 TPM device.

j) Other Hardware: Configure the graphics card and audio card types for the VM.

Option	Description
Graphics Card Type	Default graphics card type used when the VM starts. vga, virtio, qxl, and cirrus are supported and provide basic, high-definition, or high-performance video capabilities. The default value is vga for x86 VMs. ARM VMs support only virtio.
Total Graphics Memory	For vga and cirrus graphics card types, graphics memory is fixed at 16 MB. For qxl, graphics memory can be customized.
Audio Card Type	Default audio card type used when the VM starts. HDA (ICH6), HDA (ICH9), and AC97 are supported. The default value is HDA (ICH6).
Motherboard Type	Default motherboard type used when the VM starts. i440fx and q35 are supported. The default value is i440fx.

5. Configure the VM advanced settings as needed:

a) Configure general options.

Option	Description
Tag	Attach tags to identify VMs.

Option	Description
Hostname	VM hostname.
VM Scheduling Group	Join a VM scheduling group and allocate hosts according to the scheduling policy associated with the group.
Sync with Host BIOS	Whether to synchronize the Windows VM BIOS clock with the host BIOS clock. This option is disabled by default. For more information, see *VM Time Synchronization*.
User Data	User-defined data. Upload custom parameters or scripts to customize the VM or complete specific tasks. For more information, see *Import User Data to a Linux VM* and *Import User Data to a Windows VM*.

b) Configure remote access settings.

Option	Description
Console Mode	VM console mode. Options include VNC, SPICE, and VNC+SPICE. The default value is VNC.
Console Password	VM console password. You can enter a password manually or generate one randomly. Only letters, numbers, and the following special characters are supported: -`=[];',./~!@#$%^&*()_+|{}:"<>? In security settings, you can use the VNC Console Password switch to specify whether a console password is required and customize password strength. If enabled, the console password must meet the specified strength requirements.
USB Redirection	Redirect USB devices from the VDI client to the VM. This option is disabled by default.

c) Configure login authentication settings.

Option	Description
None	Do not set a login password or SSH key.
Password	System login password for the root user of a Linux VM or the Administrator user of a Windows VM. You can enter a password manually or generate one randomly. Only letters, numbers, and the

Option	Description
	following special characters are supported: -`=[];',./~!@#$%^&*()_+|{}:"<>?  **Note:** Before setting the password, make sure the VM image has `cloud-init`.
SSH Key	Select an SSH key to inject for password-free VM login. **Note:** Before injecting the SSH key, make sure the VM image has `cloud-init`.

d) Configure VMTools.

Option	Description
Failure Response Policy	Set the automatic response action after a VM failure, such as a Windows blue screen or Linux guest hang. Options include No Action, Reboot, and Shut Down.
Time Synchronization	Whether to automatically synchronize the VM base time with the host system time. This option is enabled by default.

e) Configure boot options.

Option	Description
Boot Order	Priority sequence for loading the operating system during VM boot. Hard disk, CD/DVD drive, and network boot are supported. By default, BIOS boots from the hard disk. If no boot device is found, the system cannot be loaded.
BIOS Mode	BIOS boot mode. Legacy and UEFI are supported. The default value is Legacy for x86 clusters and UEFI for ARM clusters.
Secure Boot	This option can be enabled only when BIOS Mode is UEFI. When enabled, the system validates the digital signatures of all boot

Option	Description
	components, such as the boot loader and operating system kernel, during VM startup. Only trusted software with valid signatures is loaded.
BIOS Post Delay	Duration for which the BIOS screen stays open. The default value is 10 seconds.

f) Configure other options.

Option	Description
Hide KVM Virtualization Flag	Controls whether the VM CPU virtualization flag is hidden. When enabled, `<hidden state="on">` is inserted into the `<kvm>` element in the VM definition XML. This option is disabled by default.
VMware I/O Port Simulation	Whether to allow a KVM VM to emulate I/O ports in a VMware virtualization environment so that the KVM VM can use VMware I/O port standards. This option is disabled by default. The main purposes of this option are: • Migration and compatibility: Allows VMs to be migrated from a VMware environment to a KVM environment, or allows VMware and KVM VMs to run in a mixed environment without large-scale configuration changes. • Testing and development: Developers and testers can use KVM VMs to simulate a VMware environment and test application or configuration behavior without actual VMware licenses.
Anti-Spoofing Mode	Whether to enable anti-IP/MAC spoofing and ARP spoofing protection. When enabled, the VM can communicate with external networks only by using the IP/MAC address assigned by the platform. This option is disabled by default.
Cross-Cluster HA Policy	Whether automatic VM HA migration supports cross-cluster migration. **Note:** This policy affects only automatic VM migration in shared storage scenarios, such as distributed storage, NFS storage, and SAN storage. Examples include starting a VM on another host for VM HA

Option	Description
	or migrating a VM when a host enters maintenance mode. Other operations, such as manual hot migration of a VM to another host, starting a VM on a specified host, or changing hosts based on a dynamic resource scheduling (DRS) policy, are not affected.
Hyper-V	Whether to enable Hyper-V emulation for the VM. This option is mainly used for nested virtualization scenarios on Windows VMs. It is disabled by default.
EmulatorPin	Whether to bind VM threads other than vCPU and I/O threads to host pCPUs. Binding by NUMA node is supported. By default, no binding is configured, and VM-related threads run on pCPUs according to system scheduling.
Auto-Converge	Whether to enable auto-convergence mode for KVM VM hot migration. This option is disabled by default. If the VM remains under high workload for a long time and the applicatio n is moderately sensitive to performance, enable auto-convergence mode to improve the migration success rate.
PCI Hot Plug	Whether to allow hot plugging of PCI devices from the VM. This option is enabled by default. **Note:** If a hardware compatibility error occurs during hot plug, or if the hardware device is not supported, disable this option.
CPU Vendor ID	If the host where the VM runs uses a Hygon CPU, set the VM CPU vendor ID to AuthenticAMD to improve operating system compatibility and ensure that the VM runs properly. If the value is None, some VM operating systems might be incompatible. **Note:** This parameter is hidden when the VM location is a cluster or automatically allocated.

6. Review the configuration and click **OK**.

What's next

Some VM settings depend on VMTools. After the VM is created, install VMTools so that these settings can take effect. For more information about VMTools, see *Install and Use VMTools*.

5.3.2 Import a Virtual Machine

Upload an OVF file to quickly import a VM and support migration from third-party virtualization platforms.

Prerequisites

- The platform provides VM-related system parameters that control default platform behavior. Before creating a VM, you can customize VM-related system parameters to control default VM characteristics.
- To enable VM or disk encryption, make sure that an available default key provider is configured. If you use the Native Key Provider, back it up first. For more information, see *Encryption Keys and How Encryption Works*.

Procedure

1. In the navigation pane, right-click the target cluster, host, or image, and then click **New Virtual Machine**.
2. In the **Select VM Creation Type** dialog, select **Import VM**, and then click **Next**.
3. In the **Import VM** dialog, configure the following options:

   Upload Information

Option	Description
Image Storage	Temporary storage for the uploaded template file. Standalone image storage and distributed image storage are supported. The template file is automatically deleted after the VM is generated.
Template Type	OVF template files are supported.
OVF File	Upload an OVF file. Only one file can be uploaded.
VMDK File	Upload VMDK files from the OVF template. The files must be consistent with the file configuration defined in the OVF file, including the number of files and the configuration content.

Option	Description
MF File	Upload the MF file from the OVF template. Only one file can be uploaded.

Basic Information

Option	Description
Name	VM name.
Group	VM group. If not specified, the default group is used.
Location	Host or cluster where the VM runs.
Storage Location	Data storage location of the VM disk. Local storage, NFS storage, SAN storage, and distributed storage are supported. If not specified, automatic allocation is used.
HA	Automatic restart mechanism after an abnormal VM shutdown. This option is enabled by default.
VM Encryption	Encrypts runtime data written to disk, such as memory and GPU state, by using a key from the default key provider. Enabling this option automatically enables **Disk Encryption**. For more information, see *VM Encryption*.
Disk Encryption	Encrypts all data stored on the imported VM disks by using a key from the default key provider. You can encrypt the disks even when VM encryption is disabled.
Power Status	Whether to automatically power on the VM after it is created. By default, the VM is powered on automatically.

4. Review the configuration and click **OK**.

What's next

Some VM settings depend on VMTools. After the VM is created, install VMTools so that these settings can take effect. For more information about VMTools, see *Install and Use VMTools*.

5.3.3 Register a Virtual Machine

Register VM configuration and data files to create a corresponding virtual machine record on the platform. After registration, the virtual machine can be managed by the platform.

Prerequisites

- Make sure an Advanced Edition license is installed on the platform and is valid.
- Only SAN storage, local storage, and NFS storage support VM registration. Before registration , the data storage must be attached to the corresponding cluster, and the cluster must contain hosts.
- The same VM configuration file cannot be registered multiple times on the same host.
- For a VM with TPM, make sure the corresponding key provider is configured on the platform before registration.

Procedure

1. In the navigation pane, click **Inventory** > **Data Storage**.
2. Select the target data storage, and then click **Actions** > **Register VM**.
3. In the **Register VM** dialog, configure the following options:

Option	Description
Configuration File	Select the configuration file of the VM to register.
VM Name	Set the VM name.
Group	Select the VM group.
Location	Specify where the VM runs.

4. Click **OK**.

5.3.4 Clone a Virtual Machine

To reuse the configuration of an existing virtual machine, clone the VM. NexaSphere provides full-VM cloning and supports Full Clone and Instant Full Clone.

Prerequisites

- The source VM must be running, paused, or shut down.



Note:

When cloning a running VM, only data that has been written to disk when the clone task starts is cloned. Real-time cached data is not included.

- To ensure data integrity, pause or shut down high-I/O VMs before cloning them.
- Detach shared disks from the source VM. VMs with shared disks cannot be cloned.
- The cluster where the VM resides has sufficient compute, storage, and network resources.
- If the source VM is encrypted or you plan to encrypt the cloned VM, make sure that the default key provider is available.

Procedure

1. On the target VM page, click **Actions** > **Clone** > **Clone Virtual Machine**.
2. In the **Clone Configuration** section of the **Clone Virtual Machine** dialog, select a clone method.

Option	Description
Source VM	Displays the source VM. This value cannot be changed.
Clone Method	Select a clone method: • **Full Clone**: Copies all data from the source VM. The cloned VM is always independent of the source VM. This method takes more time and storage space, but lets you specify the storage location for each disk. • **Instant Full Clone**: Quickly creates a cloned VM from the source VM, and then merges snapshots in the background so that the cloned VM becomes independent. This method takes less time and storage space. Disk storage locations are assigned automatically.  **Note:** During the snapshot merge, other VM or disk operations are performed after the merge is complete.

3. In the **Basic Info** section, configure the cloned VM.

Option	Description
Name	Specify a name for the cloned VM.

Option	Description
Quantity	Specify the number of VMs to clone. For a batch clone, the system appends sequential suffixes such as "-1", "-2", and "-3" to the VM names.
VM Encryption	Controls whether VM encryption is enabled for the cloned VM. If VM encryption is enabled for the source VM, it remains enabled for the cloned VM. Otherwise, enable this option as needed. Enabling VM encryption enables disk encryption for all disks by default. You can still adjust the encryption state of disks that were not encrypted in the source VM.
Power Status	Select **Power on after clone** to automatically start the VMs after cloning.

4. On the **Hardware Info** tab, configure the hardware for the cloned VM.

Hardware	Description
Disk	For a Full Clone, you can specify the **Storage Location** for each disk. For an Instant Full Clone, storage locations are assigned automatically. Disk encryption is handled as follows: • A disk encrypted in the source VM remains encrypted. • A disk not encrypted in the source VM remains unencrypted by default. Enable **Disk Encryption** as needed. • When **VM Encryption** is enabled, **Disk Encryption** is enabled by default for disks that were not encrypted. **Note:** • If Disk 1 of the source VM uses local storage and the other disks use different data storage, the clone fails when the local storage is unavailable. • For SAN storage, you can specify a provisioning method for the disk.

Hardware	Description
	• For nSDS distributed storage, you can select a storage pool for the disk.
NIC	The NIC configuration of the source VM is used by default. You can add or remove NICs and configure the state, NIC model, port group, number of NIC queues, MAC address, IP address, DNS, security groups, and NIC QoS. For more information about NIC options, see *Create a Virtual Machine*. **Note:** Make sure that the MAC and IP addresses of cloned VMs do not conflict with existing resources. For a batch clone, you cannot apply the same network values that must remain unique to every cloned VM.
vTPM	Displayed only if a vTPM is attached to the source VM. For **TPM Configuration**, select one of the following options: • **Retain**: Retains the TPM information of the source VM for the cloned VM. • **Reset**: Uses the default key provider to create a new TPM for the cloned VM.

5. On the **Advanced Settings** tab, select **General Options** and configure the following settings.

Option	Description
Hostname	Specify the hostname for the cloned VM. VMTools must be installed on the source VM for this setting to take effect. For a batch clone, the system appends a sequential numeric suffix to each hostname.
VM Scheduling Group	Add the cloned VM to a VM scheduling group so that a host is assigned according to the scheduling policy associated with the group. A VM can belong to only one VM scheduling group.

6. Review the configuration and click **OK**.

What's next

The clone operation copies the VM configuration, installed applications, credentials, and similar data to the cloned VM. It does not automatically copy all associated settings. Configure additional settings for the cloned VM as needed. For example, you can manually assign the same tags as the source VM.

If the source VM has a console password configured, restart the cloned VM to make the password take effect.

5.3.5 Delete a Virtual Machine

Prerequisites

- VMs in the running, paused, crashed, or unknown state cannot be moved to the recycle bin or deleted. Shut down the VM before this operation.
- (Optional) The platform provides deletion protection for VMs and disks. You can control deletion behavior by adjusting deletion policies in system parameters. By default, the VM deletion policy is delayed deletion with a 7-day retention period, and the disk deletion policy is delayed deletion with a 3-day retention period. Deleted resources are first moved to the recycle bin and are permanently deleted after the corresponding retention period.

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. Select the target VM.
3. On the VM details page, click **Actions** > **Move to Recycle Bin**.

   To delete multiple VMs, go to the Virtual Machine tab of a parent resource, such as a host, cluster, or data center. For example, on the **Virtual Machine** tab of a host, select the VMs to delete, and then click **Actions** > **Move to Recycle Bin**.

   The delete button name changes based on **VM Deletion Policy**. When the deletion policy is Immediate Deletion, the button is named "Delete". When the deletion policy is Delayed Deletion or Never Delete, the button is named "Move to Recycle Bin".
4. In the confirmation dialog, select whether to delete disks attached to the VM at the same time, excluding shared disks.

   If selected, disks are handled according to the behavior defined in the disk deletion policy.
5. After acknowledging the risk, click **OK**.

Result

After deletion, CPU, memory, IP addresses, and other resources occupied by the original VM are released immediately.

5.4 Access, Group, and Power Management

After VMs are created, you can organize them with VM groups, access them through the console or SSH, and perform power operations such as power on, shut down, reboot, and pause based on their power status.

The console proxy supports VM console access through a proxy address. Administrators can view , reconnect, or modify the console proxy address to ensure that users can access VM consoles through the proxy address.

5.4.1 VM Group Management

A VM group is a logical collection used to organize VMs based on business requirements. NexaSphere provides a default group. When you create a VM without creating or selecting a group, the VM is added to the default group. This section describes how to create and use VM groups.

Create VM Groups and Subgroups

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. Select the target data center.
3. Right-click the target data center and click **New VM Group**.
4. In the **New VM Group** dialog, configure the following options:

Option	Description
Name	Enter a name for the VM group.
Group Type	Select a group type. Options include New Group and Subgroup.
Upper-Level Group	Specify an upper-level group when you select Subgroup.

5. Click **OK**.

Manage Existing VM Groups

You can manage existing VM groups from the following entries:

- In the navigation pane, choose **Inventory** > **VM and Host**, switch to **Group View**, and select the target VM group.
- On the target VM page, click **Actions** > **System Configuration** > **Change Group**.

- Go to the Virtual Machine tab of a data center resource and perform batch operations.

Operation	Description
Add VMs to a Group	You can add VMs to a group in the following ways: • In the navigation pane, right-click the target VM group and click **New Virtual Machine**. The newly created VM belongs to this group by default. • In the navigation pane, select the target VM group. Then, on the right side of the page, click **New Virtual Machine**. The newly created VM belongs to this group by default. • On the target VM page, click **Actions** > **System Configuration** > **Change Group** to change the current VM to the target group.
Modify a VM Group Name	On the target VM group page, click **Actions** > **Edit Name**, and modify the group name.
Delete a VM Group	On the target VM group page, click **Actions** > **Delete** to delete the group. You can also delete groups in batches on the Virtual Machine tab of a data center resource. After a group is deleted, VMs in the group are moved to the default group.

5.4.2 Access a VM

After a VM is created, you can use the console to access the VM and install or configure the operating system. You can also use SSH to access a Linux VM with remote login enabled.

Console access is suitable for installing the operating system, viewing the boot process, and troubleshooting network connectivity. SSH access is suitable for routine command-line operations after the VM network is reachable.

5.4.2.1 Manage VM Access and Boot Options

5.4.2.1.1 Configure Boot Options

Prerequisites

To configure boot options, make sure the VM is stopped.

Procedure

1. On the target VM page, click **Advanced Settings**.

2. On the **Advanced Settings** tab, select **Boot Options**.
3. Click **Edit** and configure the boot order, BIOS mode, and BIOS post delay as needed.

   Boot options include the following settings:

Option	Description
Boot Order	Defines the priority sequence for loading the operating system during VM startup. For example, if you set the boot order to (1) Disk 1, (2) CD/DVD Drive, and (3) Network, the VM boots as follows: **1.** The VM first attempts to boot from Disk 1. If the boot succeeds, the VM does not try the CD/DVD drive. **2.** If booting from Disk 1 fails, the VM attempts to boot from the CD/DVD drive. If the boot succeeds, the VM does not try the network. **3.** If booting from the CD/DVD drive fails, the VM attempts to boot from the network. If the system loads successfully from the network, boot succeeds. Otherwise, boot fails and the system cannot start.
BIOS Mode	Select the BIOS boot mode based on the selected image format: • Legacy: Supports x86 architecture and all operating systems. • UEFI: Required for aarch64 architecture. Windows and CentOS are supported. Windows 7/2008 must use the compatibility support module (CSM). A mismatched BIOS mode might cause the VM to malfunction: • For qcow2 or raw images, select the same BIOS mode used when the image was created. • For iso images, select a BIOS mode as needed. The system boots and installs based on the selected mode.
BIOS Post Delay	Specifies how long the BIOS screen stays open automatically. The valid range is 1 to 60 seconds. If no operation is performed within

Option	Description
	the specified time, the system automatically proceeds from the BIOS screen to system boot.

5.4.2.1.2 Configure Remote Access

Prerequisites

To configure remote access, make sure the VM is stopped.

Procedure

1. On the target VM page, click **Advanced Settings**.
2. On the **Advanced Settings** tab, select **Remote Access**.
3. Click **Edit** and configure the console mode and console password as needed.

   Remote access includes the following settings:

Option	Description
Console Mode	Select the protocol type used to connect to the VM console: • VNC: Supports both Linux and Windows. It is mainly used for Linux server management, with typical network traffic of about 100 KB. • SPICE: Supports Linux and provides better color, audio, video , and USB capabilities. It is mainly used for virtual desktop applications, with typical network traffic of about 10 MB to 20 MB. In this mode, you can configure the number of connected displays and the video streaming mode, including off, all, and filter. • VNC+SPICE: Supports both VNC and SPICE protocols. In this mode, you can configure the number of connected displays and the video streaming mode, including off, all, and filter.
Console Password	Sets the console password. • Character requirements: Only letters, numbers, and the following special characters are supported: -`=[];',./~!@#$ %^&*()_+|{}:"<>? • Length requirements: 6 to 8 characters.

Option	Description
	The VNC console supports password strength settings.

5.4.2.1.3 Configure Login Authentication

Prerequisites

- Before setting an SSH public key, install VMTools (cloud-init). Recommended versions are 0.7.9, 17.1, 19.4, or later.
- Before changing a VM password, make sure the VM is running and VMTools (QGA) is installed .

Procedure

1. On the target VM page, click **Advanced Settings**.
2. On the **Advanced Settings** tab, select **Login Authentication**.
3. Click **Set SSH Key** or **Change VM Password**.

   Login authentication includes the following settings:

Option	Description
Set SSH Key	1. After cloud-init is installed, SSH password authentication is disabled by default. Set `ssh_pwauth` to **1** in */etc/cloud/cloud.cfg* to enable SSH password authentication. 2. Run the `ssh-keygen` command to generate an SSH key. The public key is stored in */root/.ssh/id_rsa.pub* by default. 3. Copy and paste the file content into the SSH Key input field. **Note:** • When an SSH key is injected during VM creation, it takes effect after the first boot. • When an SSH key is injected into an existing VM for the first time, restart the VM for the key to take effect. • To reinject an SSH key into an existing VM that already has an injected SSH key, run `rm -rf /var/lib/cloud/instances` to clear previous configuration, inject the new SSH key, and restart the VM for the key to take effect.

Option	Description
	• Deleting an SSH key only deletes the SSH key record from the system. It does not delete SSH key information already injected into the VM configuration. To completely delete the SSH key information, manually clean up `/root/.ssh/authorized_keys`.
Change VM Password	1. Modify the login name option, including system default and custom login name. The default login name is root for Linux and Administrator for Windows. 2. Set a new password. • Character requirements: Only letters, numbers, and the following special characters are supported: -`=[];',./~!@#$%^&*()_+|{}:"<>? VM password strength settings are supported.

5.4.2.2 Access a VM by Using the Console

Use the console to quickly access a VM and perform operations such as installing the operating system, configuring the operating system, running commands, and starting applications.

Prerequisites

- The VM must be running.
- (Optional) Configure the console mode and console password as needed. For more information, see *Configure Remote Access*.
- (Optional) Set a proxy address to access the VM console through the proxy address. For more information, see *Console Proxy Management*.

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. Select the target VM and click **Launch Console**.
   The console opens in a new browser tab.
3. Click anywhere inside the console window to start using your mouse, keyboard, and other input devices in the console.

4. Quick-access buttons are available in the left sidebar of the console window.
   - **Local Command Paster**: Click this button to open the paste command dialog. Paste the command to run into the dialog, and then click **OK** to run it in the console.
   - **Tools**: Provides hotkeys, including Ctrl, Alt, Win, Tab, Esc, and Ctrl+Alt+Del.
   - **Power Management**: Allows you to manage VM power status directly in the console. Actions include Shut Down, Reboot, Resume, Pause, and Power Off.
   - **Settings**: Allows you to enable or disable read-only mode. When read-only mode is enabled, you cannot run commands or perform other operations in the console window.

5.4.2.3 Access a VM by Using SSH

Use SSH to access a VM.

Prerequisites

- Your local computer, as the login client, has network connectivity and root access to the target VM.
- (Optional) Configure SSH key-based password-free login and the VM password as needed. For more information, see *Configure Login Authentication*.

Procedure

1. On your local computer, use remote login software to log in to the VM directly.

```
# ssh root@192.0.2.1
root@192.0.2.1's password:
Last login: Mon Sep 24 12:05:36 2021
root~#
```

2. To log in to the VM by using an SSH private key, first add the SSH public key to the target VM, and then run the following command from the login client:

```
# ssh -i ${private_key.pem} ${UserName}@${IpAddress}
```

   Here, `${private_key.pem}` is the private key file path, `${UserName}` is the login user name, and `${IpAddress}` is the IP address of the target VM.

5.4.2.4 Console Proxy Management

The console proxy is used to access VM consoles through a proxy address. By default, the proxy address is the management node IP address. When the management node is deployed in a NAT , public network access, or domain-name access scenario, modify the console proxy address and port according to the actual access path.

If a VM console fails to open, reconnect the console proxy first. After the status changes back to **Connected**, open the VM console again.

Common Entry

In the navigation pane, choose **System Management** > **Console Proxy** to go to the Console Proxy page.

Common Operations

Operation	Description
View Proxy Address	View the console proxy address. The default proxy address is the IP address of the management node.
Reconnect Console Proxy	If a VM console fails to open, click **Reconnect** on the Console Proxy page. After the status changes to **Connected**, you can open the console normally.
Modify Proxy Address	Click **Modify Proxy Address** to modify the address and port as needed. You can enter the public IP address of the management node, a NAT address, or a domain name as the console proxy address. The change takes effect immediately and does not require a management node restart.

5.4.3 VM Power Management

A VM can be in the following power statuses: Running, Powering off, Shut Down, Pausing, Paused, Powering on, Rebooting, Resuming, Unknown, and Crashed.

- Running, Shut Down, Paused, Unknown, and Crashed are final power statuses.
- Pausing, Powering off, Powering on, Rebooting, and Resuming are transitional power statuses.

Final power statuses:

- Running, Shut Down, and Paused: These statuses change based on power management operations, including Power On, Shut Down, Reboot, Resume, Pause, and Power Off.
    - If the original status is Running, Shut Down, Reboot, Pause, and Power Off are supported. If the operation succeeds, the final status changes to Shut Down, Running, Paused, and Shut Down, respectively.

  - If the original status is Shut Down, only Power On is supported. If the operation succeeds, the final status changes to Running.
  - If the original status is Paused, Shut Down, Resume, and Power Off are supported. If the operation succeeds, the final status changes to Shut Down, Running, and Shut Down, respectively.
- Unknown: When the management node cannot obtain the VM status from the host, for example when the host management network is disconnected, the VM status becomes Unknown regardless of its previous status. In the Unknown status, only Force Shut Down is supported.
- Crashed: When a running Windows VM has a blue screen or a Linux VM hangs, the VM status changes to Crashed. Based on the failure policy, the VM status changes as follows:
  - If the failure policy is No Action, the VM remains Crashed.
  - If the failure policy is Shut Down, the final VM status changes to Shut Down if shutdown succeeds. If shutdown fails, the VM remains Crashed.
  - If the failure policy is Reboot, the final VM status changes to Running if reboot succeeds. If reboot fails, the VM remains Crashed.

Transitional power statuses: Pausing, Powering off, Rebooting, and Resuming are intermediate statuses when Pause, Shut Down/Power Off/Force Shut Down, Reboot, and Resume operations change a VM from its original status to a final status. Power operations are not supported while a VM is in a transitional status.

VM Operations and Power Statuses

VM operations depend on the current power status. The following table lists the power status required before each VM operation:

Operation	VM Status
Create a VM	/
Edit a VM	Running/Shut Down/Paused
Change Group	Running/Shut Down/Paused
Power On a VM	Shut Down
Shut Down a VM	Running
Reboot a VM	Running
Resume a VM	Paused
Pause a VM	Running

Operation	VM Status
Force Shut Down a VM	Unknown
Power Off	Running
Launch Console	Running
Clone a VM	Running/Shut Down/Paused
Modify Hostname	Running/Shut Down/Paused
Bind Tag	Running/Shut Down/Paused
Unbind Tag	Running/Shut Down/Paused
Change Host	Running/Shut Down
Change Data Storage	Running/Shut Down. Supported states depend on the migration scope and data storage type. For more information, see *Change Data Storage*.
Change Host and Data Storage	Running/Shut Down
Modify CPU Hot Plug and Memory Hot Plug	Running/Shut Down
Configure GPU Device	Shut Down
Expand System Disk	Running/Shut Down
Change Owner	Running/Shut Down
Change System	Shut Down
Reset System	Shut Down
Configure Boot Order	Running/Shut Down
Configure HA	Running/Shut Down
Configure BIOS Clock Synchronization	Running/Shut Down/Paused
Configure Time Synchronization	Running/Shut Down/Paused
Configure Failure Policy	Running
Set SSH Key	Running/Shut Down
Change VM Password	Running
Set Console Password	Running/Shut Down
Switch Console Mode	Running/Shut Down/Paused
Configure CPU and Memory Resource Priority	Running/Shut Down/Paused
Cross-Cluster HA Policy	Running/Shut Down/Paused

Operation	VM Status
USB Redirection	Running/Shut Down/Paused
Create Snapshot	Running/Shut Down/Paused
Attach Disk	Running/Shut Down
Detach Disk	Running/Shut Down
Create Image	Running/Shut Down
Attach ISO	Running/Shut Down
Detach ISO	Running/Shut Down
Join VM Scheduling Group	Running/Shut Down
Remove from VM Scheduling Group	Running/Shut Down
Export VM	Shut Down
Delete VM (Move to Recycle Bin)	Shut Down
Permanently Delete/Restore VM	Deleted
Copy URL	Running/Shut Down/Paused
Download Exported VM	Running/Shut Down/Paused
Delete Export Record	Running/Shut Down/Paused

5.5 Install and Use VMTools

VMTools is a set of drivers and tools installed in VMs to enhance NexaSphere management of guest operating systems. After VMTools is installed, the platform can read internal VM information and support configuration delivery, configuration retrieval, time synchronization, and advanced monitoring.

VMTools Components

VMTools primarily includes the following tools and drivers. The components vary by operating system.

OS	VMTools Component	Description
Linux	Advanced Monitoring Agent	An agent installed in a VM to collect advanced monitoring data, including CPU, memory, and disk capacity data.
	QEMU Guest Agent (QGA)	An application that enables interaction between VMs and hosts without relying on the network. QGA can retrieve and deliver VM configurations and push monitoring data.

OS	VMTools Component	Description
		Note: Linux operating systems usually include virtio drivers and cloud-init by default. Therefore, VMTools for Linux does not include these two tools. If your Linux operating system does not include them, install them separately.
Windows	Advanced Monitoring Agent	An agent installed in a VM to collect advanced monitoring data, including CPU, memory, and disk capacity data.
	QGA	An application that enables interaction between VMs and hosts without relying on the network. QGA can retrieve and deliver VM configurations and push monitoring data.
	Virtio Driver	A set of drivers that improves VM performance. • SCSI controller driver: Improves VM disk performance in a virtualization environment. • PCI simple communications controller driver: Enables communication between a VM and the underlying KVM virtualization layer. • PCI device driver: Enables balloon memory scaling. • Ethernet adapter driver: Improves VM network performance in a virtualization environment.
	Cloudbase-init	A component that supports User Data import and other customization features.
FreeBSD	Advanced Monitoring Agent	An agent installed in a VM to collect advanced monitoring data, including CPU, memory, and disk capacity data.

VMTools Use Cases

After VMTools is installed on a VM, the platform can support advanced monitoring, configuration delivery, and configuration retrieval.

Advanced monitoring	The Advanced Monitoring Agent collects monitoring data from inside the VM and periodically pushes the data to the host through DHCP or QGA. You can use line charts to view real-time trends in VM performance, including CPU, memory, and disk capacity.
Configuration delivery	QGA in VMTools delivers parameters configured on the platform to the guest operating system, including the VM hostname and network configurations.

- Hostname: After the latest VMTools is installed, a new hostname configured for a running VM is delivered through QGA and takes effect without restarting the VM.

  Note:
  To deliver a hostname through QGA, install VMTools before you modify the hostname.

- Network configuration: If DHCP is not enabled for a distributed port group, parameters configured on the platform, such as the IP address, DNS, and MTU , must be delivered to the NIC through QGA. After VMTools is installed, the specified IP address and other parameters automatically take effect.

Configuration retrieval

QGA in VMTools automatically retrieves the IP address configured for a VM NIC. The retrieved IP address can be displayed and managed on the platform.

Note:

- If DHCP is not enabled for the distributed port group to which the VM NIC belongs, and VMTools is installed on the running VM, the platform can automatically retrieve the VM NIC configuration.
- If an IP address is already configured inside the VM, the retrieved IP address overwrites the IP address configured on the platform.
- If the VM NIC IP address conflicts with an IP address in the port group, an alert is generated.

5.5.1 Install VMTools on a Linux VM

The installation process mounts the VMTools ISO first. You then open the VM console and run the installation commands manually.

Prerequisites

- Make sure the VM is running.
- Make sure Linux command-line tools are installed on the VM, such as `tar`, `wget`, and `curl`.

Procedure

1. On the target VM page, click **Actions** > **VMTools** > **Install VMTools**.
2. In the **Install VMTools** dialog, click **Next: Install on VM Console**.

When mounting the VMTools ISO, the system selects a CD/DVD drive according to the following rules:

- If an empty CD/DVD drive exists, the system automatically selects the first available empty CD/DVD drive and mounts the VMTools ISO.
- If no empty CD/DVD drive exists, the system selects CD/DVD Drive 1, automatically unmounts the original image, and mounts the VMTools ISO.

3. Click **Copy Commands and Launch Console**.

   After you click the button, the following commands are copied and the VM console opens automatically:

```
# Create a mount point directory
mkdir /mnt/cdrom
# Mount the CD-ROM image
mount /dev/cdrom /mnt/cdrom
# Install VMTools
cd /mnt/cdrom/
bash ./zs-tools-install.sh
# Unmount the CD-ROM image (optional)
cd ~
umount /mnt/cdrom
```

4. In the VM console, click **Local Command Paster**.
5. In the **Paste Command** dialog, paste the commands copied in the previous step, and then click **OK** to run the commands.

What's next

To update VMTools, on the target VM page, click **Actions** > **VMTools** > **Reinstall VMTools**.

5.5.2 Install VMTools on a Windows VM

The installation process mounts the VMTools ISO first. You then open the VM console and install VMTools as prompted.

Prerequisites

Make sure the VM is running.

Procedure

1. On the target VM page, click **Actions** > **VMTools** > **Install VMTools**.
2. In the **Install VMTools** dialog, click **Next: Install on VM Console**.

   When mounting the VMTools ISO, the system selects a CD/DVD drive according to the following rules:

- If an empty CD/DVD drive exists, the system automatically selects the first available empty CD/DVD drive and mounts the VMTools ISO.
- If no empty CD/DVD drive exists, the system selects CD/DVD Drive 1, automatically unmounts the original image, and mounts the VMTools ISO.

3. Click **Launch Console**.
4. Open the VM console and install VMTools as prompted.
    1. Mount the VMTools image: In the VMTools installation prompt, click Install Now to mount the VMTools image to the virtual CD/DVD drive.
    2. Install VMTools: Run the VMTools installer to install VMTools, common tools, and VirtIO drivers in sequence.
    3. Complete the installation and restart the VM: After the installation is complete, restart the VM for the configuration to take effect.

     **Note:**
    Windows Server 2008 R2 does not support silent installation of VirtIO drivers. If Windows displays a driver publisher or security prompt during installation, manually trust the driver publisher and continue.

What's next

To update VMTools, on the target VM page, click **Actions** > **VMTools** > **Reinstall VMTools**.

5.5.3 VMTools Supported OS List

VMTools support varies by operating system. The following tables list VMTools support for each operating system.

 **Note:**
In the tables, **Yes** means that the component can be installed through VMTools and related features can be used normally. **No** means that the component cannot be installed through VMTools, or features that depend on the component might be unavailable after installation. It does not mean the component itself is unsupported. You can still install the component by other means.

Linux Operating Systems

Operating System	Version	Advanced Monitoring Agent	QGA
CentOS	CentOS 6.5 64-bit	No	Yes
	CentOS 6.8 64-bit	Yes	Yes
	CentOS 6.9 64-bit	Yes	Yes
	CentOS 6.10 64-bit	Yes	No
	CentOS 7.2 64-bit	Yes	Yes
	CentOS 7.3 64-bit	Yes	Yes
	CentOS 7.4 64-bit	Yes	Yes
	CentOS 7.5 64-bit	Yes	Yes
	CentOS 7.6 64-bit	Yes	Yes
	CentOS 7.9 64-bit	Yes	Yes
	CentOS 8.0 64-bit	Yes	Yes
RHEL	Redhat Enterprise Linux Server 6.9 64-bit	Yes	No
	Redhat Enterprise Linux Server 7.0 64-bit	Yes	No
	Redhat Enterprise Linux Server 7.1 64-bit	Yes	No
	Redhat Enterprise Linux Server 7.2 64-bit	Yes	No
	Redhat Enterprise Linux Server 7.3 64-bit	Yes	No
	Redhat Enterprise Linux Server 7.4 64-bit	Yes	Yes
	Redhat Enterprise Linux Server 7.5 64-bit	Yes	Yes
	Redhat Enterprise Linux Server 7.6 64-bit	Yes	Yes
Fedora	Fedora 30 64-bit	Yes	Yes
	Fedora 31 64-bit	Yes	Yes

Operating System	Version	Advanced Monitoring Agent	QGA
Debian	Debian 9.9 64-bit	Yes	Yes
	Debian 10.13 64-bit	Yes	Yes
	Debian 11.9 64-bit	Yes	Yes
	Debian 12.5 64-bit	Yes	Yes
Ubuntu	Ubuntu 14.04 64-bit	Yes	Yes
	Ubuntu 16.04 64-bit	Yes	Yes
	Ubuntu 16.10 64-bit	Yes	Yes
	Ubuntu 18.04 64-bit	Yes	Yes
	Ubuntu 20.04 64-bit	Yes	Yes
	Ubuntu 22.04 64-bit	Yes	Yes
	Ubuntu 24.04 64-bit	Yes	Yes
Kylin	Kylin V4.0.2 64-bit	Yes	No
	Kylin V10 SP1(0518) 64-bit	Yes	Yes
	Kylin V10 SP2 64-bit	Yes	Yes
	Kylin V10 SP3 64-bit	Yes	Yes
NeoKylin	NeoKylin V7.0 64-bit	Yes	Yes
	NeoKylin V7update6 64-bit	Yes	Yes
OpenSUSE	OpenSUSE Leap 15.0 64-bit	Yes	Yes
SLES	SUSE Linux Enterprise Server 11 64-bit	Yes	Yes
	SUSE Linux Enterprise Server 12 64-bit	Yes	Yes
	SUSE Linux Enterprise Server 15 64-bit	Yes	Yes
	SUSE Linux Enterprise Desktop 12 64-bit	Yes	Yes
	SUSE Linux Enterprise Desktop 15 64-bit	Yes	Yes
UOS	UOS V20 1050e	Yes	Yes

Operating System	Version	Advanced Monitoring Agent	QGA
Oracle Linux	Oracle Linux 7.9	Yes	Yes
OpenEuler	OpenEuler 20 64-bit	Yes	Yes
	OpenEuler 22 64-bit	Yes	Yes
Alma Linux	Alma Linux 9.3 64-bit	Yes	Yes

Windows Operating Systems

Operating System	Version	Advanced Monitoring Agent	QGA	VirtIO Driver	Cloudbase-init
Windows	Windows Server 2008 R2 64-bit	Yes	Yes	Yes	No
	Windows Server 2012 64-bit	Yes	Yes	Yes	Yes
	Windows Server 2016 64-bit	Yes	Yes	Yes	Yes
	Windows Server 2019 64-bit	Yes	Yes	Yes	Yes
	Windows Server 2022 64-bit	Yes	Yes	Yes	Yes
	Windows Server 2025 64-bit	Yes	Yes	Yes	Yes

Note:

Windows Server 2008 R2 supports installing VirtIO drivers through VMTools, but it does not support silent installation. If Windows displays a driver publisher or security prompt during installation, manually trust the driver publisher and continue.

FreeBSD Operating System

Operating System	Version	Advanced Monitoring Agent
FreeBSD	FreeBSD 11 64-bit	Yes
	FreeBSD 12 64-bit	Yes

Operating System	Version	Advanced Monitoring Agent
	FreeBSD 13 64-bit	Yes

5.6 Configure a VM

VM configuration includes basic information, hardware resources, encryption, access options, time synchronization, User Data, CPU features, QoS, and resource contention settings. Some settings depend on the VM status, guest operating system capabilities, or VMTools.

Before modifying VM configuration, confirm the maintenance window, VM status, and conditions for the settings to take effect. For settings related to storage, networking, resource scheduling, or security mechanisms, see the corresponding manuals for complete mechanisms.

5.6.1 Modify VM Encryption Settings

Enable or disable encryption for an existing VM, or change the encryption state of individual disks.

Prerequisites

- The VM is shut down.
- An available default key provider is configured. If you use the Native Key Provider, back it up first. For more information, see *Encryption Keys and How Encryption Works*.
- Shared disks and RDM disks do not support disk encryption. Before changing encryption settings, make sure that the target disk is neither a shared disk nor an RDM disk.

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. Select the target VM and click **Modify Configuration**.
3. In the **Basic Info** section of the **Modify Configuration** dialog, enable or disable **VM Encryption** as needed.
   - Enabling VM encryption encrypts runtime data written to disk, such as memory and GPU state, by using a key from the default key provider. It also automatically enables **Disk Encryption** for the existing disks.
   - Disabling VM encryption does not automatically disable **Disk Encryption**. You can keep individual disks encrypted or disable disk encryption in the next step.
4. In the **Hardware Info** section, select a disk and enable or disable **Disk Encryption** as needed.

   Disk encryption protects all data stored on the selected disk. You can encrypt an individual disk even when VM encryption is disabled.

5. Review the configuration and click **OK**.



Note:

Changing the encryption state of an existing VM or disk requires processing existing data and usually takes longer than enabling encryption during VM creation. Source data retained during conversion is displayed on the **Data Cleanup** tab of the data storage. After verifying that the converted VM works as expected, clean up the source data as needed.

5.6.2 VM Time Synchronization

Time consistency between VMs and hosts is important for workloads running on the platform. NexaSphere provides two time synchronization mechanisms: BIOS clock synchronization and VM time synchronization.

BIOS Clock Synchronization

BIOS clock synchronization is designed for Windows VMs. For a Windows VM, hardware clock time comes from the host, and system time is equivalent to hardware clock time. You can enable BIOS clock synchronization to keep the Windows VM hardware time, system time, and host hardware time consistent.

When creating a VM, enable **Sync with Host BIOS Time** in **Advanced Settings** > **General Options**. After it is enabled, the VM periodically synchronizes the host BIOS clock automatically.

Time Synchronization

Time synchronization is designed to synchronize the system time of Windows and Linux VMs with the host system time. For a Linux VM, hardware time comes from the host, while system time is calculated separately. System time is synchronized with hardware time only at specific moments, such as when the system starts. At other times, system time is independent of hardware time.

VM time synchronization mechanism:

- Network: Unlike NTP time synchronization, this mechanism does not depend on a specific external network. It involves only communication between the host and the VM.
- Time zone: The VM periodically synchronizes the time zone of the host where it runs to keep them consistent.
- Synchronization interval: By default, synchronization is performed every 60 seconds. You can modify the synchronization interval by using the CLI. Supported values are 60 seconds (1 minute), 600 seconds (10 minutes), 1800 seconds (30 minutes), 3600 seconds (1 hour), 7200

seconds (2 hours), 21600 seconds (6 hours), 43200 seconds (12 hours), and 86400 seconds (1 day).

```
# Modify the synchronization interval by using the CLI
[root@localhost ~]# UpdateResourceConfig vm=<vm_uuid> category=vm \
name=vm.clock.sync.interval.in.seconds value=<intervalInSeconds>

# vm_uuid is the VM UUID
# intervalInSeconds is the synchronization interval to set, in
seconds
```

- Synchronization policy: The current policy is immediate consistency. Regardless of whether the VM time is ahead of or behind the host time, when synchronization is triggered, the VM time immediately becomes consistent with the host time.

To configure time synchronization for a single VM:

1. Make sure QEMU Guest Agent (QGA) is installed on the VM and is running. For more information, see *Install and Use VMTools*.
2. Disable other time source synchronization policies on the VM. This is recommended.
3. On the VM page, click **Actions** > **Advanced Settings** > **Modify VMTools** and enable time synchronization. After it is enabled, the host system time is synchronized automatically every 60 seconds.

5.6.3 Import User Data

5.6.3.1 Import User Data to a Linux VM

User Data is user-defined data. You can use custom parameters or scripts to customize a Linux VM or perform specific tasks when creating the VM.

Prerequisites

- Before importing User Data, make sure cloud-init is installed in the VM image. Recommended versions are 0.7.9, 17.1, 19.4, or later.
- If you use User Data to set the hostname and SSH login password, do not set the hostname and password again on the platform to avoid conflicts.
- When creating a Linux VM from an image that has cloud-init installed, you must import User Data. Otherwise, the cloud-init task waits until timeout.

Procedure

1. Prepare the User Data script for the required functions.

Script example:

```
#cloud-config
users:
 - name: test
   shell: /bin/bash
   groups: users
   sudo: ['ALL=(ALL) NOPASSWD:ALL']
   ssh-authorized-keys:
       - ssh-rsa AAAAB3NzaC1LXCJfjroD1lT root@10-0-0-18
bootcmd:
 - mkdir /tmp/temp
write_files:
 - path: /tmp/Cloud_config
   content: |
       Hello,world!
   permissions: '0755'
fqdn: Perf-test
disable_root: false
ssh_pwauth: yes
chpasswd:
  list: |
      root:word
  expire: False
runcmd:
 - echo ls -l / >/root/list.sh
```

The preceding example implements the following functions:

1. Creates the user test and uses an SSH key when the VM is created.
2. Writes the file */etc/hosts*, creates the */tmp/temp* directory, and creates a file with content during startup.
3. Sets the hostname, enables the root user, allows SSH password login, and changes the root password.
4. Runs the `echo ls -l /` command.

2. In the navigation pane, select **Inventory**.
3. In the inventory, right-click a cluster, host, or image, and then click **New Virtual Machine**.
4. In the **New Virtual Machine** dialog, select **Advanced Settings** > **General Options**.
5. Paste the script into the **User Data** input box.
6. Complete the VM creation process to import User Data.

5.6.3.2 Import User Data to a Windows VM

User Data is user-defined data. You can use custom parameters or scripts to customize a Windows VM or perform specific tasks when creating the VM.

Prerequisites

- Before importing User Data, make sure Cloudbase-init is installed in the VM image. The recommended version is 0.9.11. For more information about Cloudbase-init, see *Cloudbase official documentation*.
- If you use User Data to set the hostname and SSH login password, do not set the hostname and password again on the platform to avoid conflicts.
- When creating a Windows VM from an image that has Cloudbase-init installed, you must import User Data. Otherwise, the Cloudbase-init task waits until timeout.

Procedure

1. Prepare the User Data script for the required functions.

   Script example:

   ```
   #cloud-config
   write_files:
      -   encoding: b64
          content: NDI=
          path: C:\b64
          permissions: '0644'
      -   encoding: base64
          content: NDI=
          path: C:\b64_1
          permissions: '0644'
      -   encoding: gzip
          content: !!binary |
              H4sIAGUfoFQC/zMxAgCIsCQyAgAAAA==
          path: C:\gzip
          permissions: '0644'
   ```

   The preceding example creates the following three files on drive C during VM startup: *b64*, *b64_1*, and *gzip*.

2. In the navigation pane, select **Inventory**.
3. In the inventory, right-click a cluster, host, or image, and then click **New Virtual Machine**.
4. In the **New Virtual Machine** dialog, select **Advanced Settings** > **General Options**.
5. Paste the script into the **User Data** input box.
6. Complete the VM creation process to import User Data.

5.6.4 Configure VM Resource Contention

NexaSphere provides the following settings to improve VM resource contention capability:

- *CPU NUMA Binding*
- *EmulatorPin*
- *CPU and Memory Resource Priority*

CPU NUMA Binding

CPU NUMA binding: CPU pinning assigns the virtual CPUs (vCPUs) of a virtual machine to specific physical CPUs (pCPUs) of the host, which improves VM performance.

Consider CPU binding in the following scenarios to improve VM performance:

- CPU-intensive small-application scenario:

  CPU-intensive applications require substantial CPU resources. If many CPU-intensive small applications run on VMs, CPU resource contention might occur. CPU binding allows applications to run on specific physical CPUs, helping reduce contention and improve system performance.
- Uneven multi-core CPU load scenario:

  If multiple applications are concentrated on one or a few CPUs, use CPU binding to manually adjust CPU load distribution. The change takes effect immediately.

To configure CPU binding:

1. On the target VM page, click **Shut Down** to shut down the VM.
2. Click **Modify Configuration** and select **Hardware** > **CPU** > **CPU NUMA Binding**, and then select Intelligent Binding or Manual Binding as needed:
   - **Intelligent Binding**: Binds VM vCPUs to pCPUs in pNUMA nodes one to one in descending order of pNUMA node ID. After all pCPUs in a pNUMA node are bound to VM vCPUs, pCPUs in the next pNUMA node are bound in sequence. If all pCPUs are bound but some vCPUs are still unbound, the system loops back to the first bound pNUMA node and continues one-to-one binding.
   - **Manual Binding**: Manually bind vCPUs to pCPUs based on the host pNUMA topology. All vCPUs must be bound.
     - A vCPU can be bound to one or more pCPUs, and a pCPU can be bound to one or more vCPUs. All pCPUs bound to one vCPU must belong to the same pNUMA node.
     - The average usage of each pCPU in the last 15 minutes is displayed to help you select the best pCPU for binding.

 **Note:**

- If CPU overcommitment is configured, the number of vCPUs might be greater than the number of pCPUs. In this case, make sure the number of vCPUs does not exceed the number of bound pCPUs. Otherwise, VM performance might be significantly affected.
- After CPU NUMA binding is configured, modifying the number of VM CPU cores automatically cancels CPU NUMA binding.

3. Click **OK** to complete CPU NUMA binding. You can click **Power On** to power on the VM again.

EmulatorPin

EmulatorPin: EmulatorPin assigns all other threads than virtual CPU (vCPU) threads and IO threads of a virtual machine to physical CPUs (pCPUs) of the host so that these threads run on assigned pCPUs.

Consider EmulatorPin in the following scenario:

- Parallel multi-service scenario:

  When multiple VMs run different workloads on one host, they consume different amounts of resources. EmulatorPin binds the VM QEMU main thread to pCPUs, isolating primary service processes across VMs and keeping system performance relatively stable.

To configure EmulatorPin binding:

1. On the target VM page, click **Actions** > **Advanced Settings** > **Modify Other Settings**, and then select **EmulatorPin**: Bind VM threads other than vCPU and I/O threads to host pCPUs based on the host pNUMA topology.
2. Click **OK** to complete EmulatorPin binding.

CPU and Memory Resource Priority

If VM A shares CPU and memory resources of Host A with other VMs, and VM A has higher business importance and priority, resource contention might occur when resources are tight, for example when host load is high. In this case, set the CPU Resource Priority and Memory Resource Priority of VM A to **High**, while other VMs use the default value **Normal**, to improve VM A resource contention capability.

To set resource priority:

1. On the target VM page, click **Modify Configuration** and select **Hardware** > **CPU** > **CPU Resource Priority**, and then select a priority as needed.
2. Click **OK** to complete resource priority configuration.

5.6.5 Configure VM CPU Features

NexaSphere provides CPU mode settings. You can use CPU mode to control whether the VM CPU model matches the host CPU model, so that the VM can inherit some or all host CPU features for specific workload requirements.

VM CPU Modes

NexaSphere supports the following CPU modes:

- None (default): The VM CPU model is simulated by QEMU and inherits a limited set of CPU features from the host. This mode is recommended for migration scenarios.
- Compatible: The VM CPU model is close to or consistent with the host CPU model. For example, both can be displayed as Haswell Intel CPU. Compared with None mode, this mode allows the VM to inherit more CPU features from the host and can be used for migration scenarios. This setting is not supported for aarch architecture clusters or VMs.
- Passthrough: The VM CPU model is consistent with the host CPU model. VM CPU features are also consistent with host CPU features, such as extended page tables, huge pages, and virtualization. Compared with None, Compatible, and Custom modes, this mode inherits the most host CPU features and is suitable for workloads with high VM feature requirements.
- Custom (a specific CPU model): The VM is configured with the specified custom CPU model. After a custom CPU model is configured, the VM might have CPU features that differ from the previous custom model.

Configure VM CPU Mode

NexaSphere provides CPU mode settings at the VM level and cluster level. VM-level settings take precedence over cluster-level settings. The following procedure uses VM-level settings as an example:

1. On the target VM page, click **Shut Down** to shut down the VM.
2. Click **Modify Configuration** and select **Hardware** > **CPU** > **CPU Mode**, and then select the required CPU mode.
3. Click **OK** to complete CPU mode configuration. You can click **Power On** to power on the VM again.

5.6.6 Configure VM QoS

QoS helps address network latency and congestion by setting I/O bandwidth thresholds. When the network is overloaded or congested, QoS helps prevent service traffic from being delayed or dropped and keeps the network running efficiently.

You can use QoS to limit disk bandwidth, disk IOPS, and NIC inbound and outbound bandwidth. If QoS is not configured, the system does not limit the corresponding I/O or network bandwidth.

Disk QoS

To configure disk QoS:

1. On the target VM page, click **Modify Configuration** and select **Hardware** > **Disk {$n}** > **QoS**, and enable the QoS switch to configure QoS for Disk n, such as Disk 1.

   Disk QoS supports bandwidth limits and IOPS limits. After QoS is enabled, configure at least one type of limit:

   - **Bandwidth Limit**: Sets the upper limit of disk read/write throughput per second. Basic units include MB/s and GB/s. Do not set this value too low to avoid VM exceptions.
     - **Total Speed**: Sets the upper limit of total disk read/write throughput.
     - **Read/Write Speed**: Sets the upper limits of read throughput and write throughput separately.
   - **IOPS Limit**: Sets the upper limit of disk read/write operations per second.
     - **Total IOPS**: Sets the upper limit of total disk read/write IOPS.
     - **Read/Write IOPS**: Sets the upper limits of read IOPS and write IOPS separately.
2. Click **OK** to complete disk QoS configuration.

Network QoS

To configure network QoS:

1. On the target VM page, click **Modify Configuration** and select **Hardware** > **NIC {$n}** > **QoS**, and enable the QoS switch to configure QoS for NIC n, such as NIC 1.

   QoS supports send bandwidth and receive bandwidth. Basic units include Kbps, Mbps, and Gbps. After QoS is enabled, configure at least one bandwidth limit:

   - **Transmit Bandwidth**: Upper limit of network bandwidth uploaded from the VM.
   - **Receive Bandwidth**: Upper limit of network bandwidth downloaded to the VM.
2. Click **OK** to complete network QoS configuration.

5.7 Migrate VMs

NexaSphere provides internal VM migration capabilities for changing the host, data storage, or both within an existing NexaSphere environment. When changing data storage, you can migrate the entire VM or only selected data disks. Depending on the VM power state, data storage type, and business continuity requirements, you can use hot migration or cold migration.

Before migration, make sure that the destination host, destination data storage, network connectivi ty, and VM state meet the migration requirements.

5.7.1 Change Host

5.7.1.1 Hot Migrate a VM to Another Host

Hot migration migrates a running virtual machine. This action mainly copies the CPU register state and memory information.

Prerequisites

- The VM resides on local storage, NFS storage, nSDS distributed storage, SAN storage, or ZHPS distributed storage.
- If the VM has data disks, ISOs, or peripheral devices attached, detach them before migration . A VM on local storage can be hot migrated with local data disks attached. A VM on shared storage can be hot migrated with data disks, shared disks, and ISOs attached.
- If vNUMA is enabled for the VM, make sure that the pNUMA topology of the destination host is consistent with that of the source host.
- If the VM has VF NICs attached, make sure that the following requirements are met before hot migration:
  - VMTools is installed on the VM and is running properly.
  - If the driver version inside the VM is too old, migration might fail. Check and update the driver version as needed.
- If the selected VM is associated with VM scheduling policies, modify the policies based on your service requirements to avoid policy conflicts.

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. Select the target VM and click **Actions** > **Migration Management** > **Change Host**.
3. In the **Change Host** dialog, configure the migration settings.

Option	Description
Destination Host	Select the host to which the VM is migrated.
Migration Encryption	Encrypts memory and runtime state data transferred during hot migration. If VM encryption is enabled, this option is enabled automatically and cannot be disabled. If VM encryption is not enabled, this option is disabled by default and can be enabled as needed. The disk encryption state does not affect this option. For more information, see *Interoperability with Other Features*.
Auto-Converge	If the VM has been running under a high service load for a long time and the service can tolerate a moderate performance impact, enable this option to improve the migration success rate.

4. Review the configuration and click **OK**.

> **Note:**
> If the selected VM has VF NICs attached, short network interruptions might occur during migration depending on the physical network environment. Proceed with caution for service-sensitive VMs.

What's next

After migration, you can re-attach the detached data disks, ISOs, or peripheral devices to the VM.

5.7.1.2 Cold Migrate a VM to Another Host

Cold migration migrates a powered-off virtual machine.

Prerequisites

- The VM resides on local storage.
- If the VM has ISOs or peripheral devices attached, detach them before migration.

> **Note:**
> A VM on local storage can be cold migrated with local data disks attached.

- If the selected VM is associated with VM scheduling policies, modify the policies based on your service requirements to avoid policy conflicts.

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.

2. Select the target VM and click **Shut Down**.
3. After the VM is shut down, click **Actions** > **Migration Management** > **Change Host**.
4. Configure the migration.

Option	Description
Destination Host	Selects the host to which the VM is migrated.
Migration Encryption	If changing the host also transfers local disk data, this option encrypts the data in transit. If VM encryption is enabled, the option is enabled automatically and cannot be disabled. Otherwise, it is disabled by default and can be enabled as needed. The disk encryption state does not affect this option. For more information, see *Interoperability with Other Features*.

5. Review the configuration and click **OK**.

What's next

After migration, you can re-attach the detached data disks, ISOs, or peripheral devices to the VM.

5.7.2 Change Data Storage

Changing data storage migrates the system disk or data disks of a VM without changing the host on which the VM runs. You can migrate the entire VM or only selected data disks.

If encryption is enabled for the VM or any disk to be migrated, see *Interoperability with Other Features*.

Migration Scope

Scope	Description
Entire VM	Migrates the system disk of the VM. You can also migrate the attached data disks as needed.
Data Disk	Migrates one or more selected data disks without changing the data storage of the system disk. You can specify a different destination data storage for each data disk.

Supported Migration Methods

Migration Scope	Source and Destination Storage	VM State
Entire VM	SAN storage to SAN storage	Running or shut down
Entire VM	Between SAN storage and nSDS distributed storage	Running
Data Disk	SAN storage to SAN storage	Running or shut down
Data Disk	NFS storage to NFS storage	Shut down
Data Disk	nSDS distributed storage to nSDS distributed storage	Shut down

The source and destination data storage must be attached to the same cluster. Data disks on local storage or other unsupported storage types cannot be migrated.

Note:
If the destination data storage contains source data retained after a failed storage migration, remove the retained data on the **Data Cleanup** tab of the destination data storage details page before retrying the migration.

5.7.2.1 Hot Migrate a VM from SAN Storage to SAN Storage

Prerequisites

- Make sure that the source data storage and destination data storage are attached to the same cluster.
- If the VM has ISOs or block devices attached, detach them before migration.
- If a shared data disk is attached to the VM, the shared disk cannot be migrated together with the VM. Detach the shared disk before migration.
- If the VM has VF NICs attached, make sure that the following requirements are met before hot migration:
  - VMTools is installed on the VM and is running properly.
  - If the driver version inside the VM is too old, migration might fail. Check and update the driver version as needed.

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. Select the target VM and click **Actions** > **Migration Management** > **Change Data Storage**.
3. In the risk confirmation dialog, review the migration risks, select **I understand the risks above**, and click **OK**.
4. In the Migration Scope section of the **Change Data Storage** dialog, select **Entire VM**.
5. Configure the migration.

Option	Description
Destination Data Storage	Selects the SAN data storage to which the system disk is migrated.
Migrate Data Disks	Specifies whether to migrate all data disks attached to the VM.
Migration Bandwidth	Limits the bandwidth used for migration. Leave this option empty to apply no bandwidth limit.
Migration Encryption	Encrypts data transferred during migration. If VM encryption is enabled, this option is enabled automatically and cannot be disabled. If VM encryption is not enabled, this option is disabled by default and can be enabled as needed. The disk encryption state does not affect this option. For more information, see *Interoperability with Other Features*.

6. Review the configuration and click **OK**.

Note:

- If the selected VM has VF NICs attached, short network interruptions might occur during migration depending on the physical network environment. Proceed with caution for service-sensitive VMs.
- If the selected VM has snapshots, storage migration while the VM is running deletes these snapshots. Proceed with caution. To keep snapshots, shut down the VM before migration.

What's next

After migration, you can re-attach the detached data disks, ISOs, or peripheral devices to the VM.

5.7.2.2 Cold Migrate a VM from SAN Storage to SAN Storage

Prerequisites

- Make sure that the source data storage and destination data storage are attached to the same cluster.
- If the VM has ISOs or block devices attached, detach them before migration.
- If a shared data disk is attached to the VM, the shared disk cannot be migrated together with the VM. Detach the shared disk before migration.

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. Select the target VM and click **Shut Down**.
3. After the VM is shut down, click **Actions** > **Migration Management** > **Change Data Storage**.
4. In the risk confirmation dialog, review the migration risks, select **I understand the risks above**, and click **OK**.
5. In the Migration Scope section of the **Change Data Storage** dialog, select **Entire VM**.
6. Configure the migration.

Option	Description
Destination Data Storage	Selects the SAN data storage to which the system disk is migrated.
Migrate Data Disks	Specifies whether to migrate all data disks attached to the VM.
Migration Bandwidth	Limits the bandwidth used for migration. Leave this option empty to apply no bandwidth limit.
Migration Encryption	Encrypts data transferred during migration. If VM encryption is enabled, this option is enabled automatically and cannot be disabled. If VM encryption is not enabled, this option is disabled by default and can be enabled as needed. The disk encryption state does not affect this option. For more information, see *Interoperability with Other Features*.

7. Review the configuration and click **OK**.

What's next

After migration, you can re-attach the detached data disks, ISOs, or peripheral devices to the VM.

5.7.2.3 Hot Migrate a VM Between SAN Storage and nSDS Distributed Storage

Prerequisites

- Make sure that the source data storage and destination data storage are attached to the same cluster.
- If the VM has ISOs or block devices attached, detach them before migration.
- If a shared data disk is attached to the VM, the shared disk cannot be migrated together with the VM. Detach the shared disk before migration.
- If the VM has VF NICs attached, make sure that the following requirements are met before migration:
  - VMTools is installed on the VM and is running properly.
  - If the driver version inside the VM is too old, migration might fail. Check and update the driver version as needed.

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. Select the target VM and click **Actions** > **Migration Management** > **Change Data Storage**.
3. In the risk confirmation dialog, review the migration risks, select **I understand the risks above**, and click **OK**.
4. In the Migration Scope section of the **Change Data Storage** dialog, select **Entire VM**.
5. Configure the migration.

Option	Description
Destination Data Storage	Selects the destination data storage to which the system disk is migrated.
Disk 1 Storage Pool	Selects the destination storage pool for the system disk when the destination is nSDS distributed storage.
Migrate Data Disks	Specifies whether to migrate all data disks attached to the VM. When the destination is nSDS distributed storage, the data disks are migrated to the storage pool selected for Disk 1.
Migration Bandwidth	Limits the bandwidth used for migration. Leave this option empty to apply no bandwidth limit.

Option	Description
Migration Encryption	Encrypts data transferred during migration. If VM encryption is enabled, this option is enabled automatically and cannot be disabled. If VM encryption is not enabled, this option is disabled by default and can be enabled as needed. The disk encryption state does not affect this option. For more information, see *Interoperability with Other Features*.

6. Review the configuration and click **OK**.



Note:

- If the selected VM has VF NICs attached, short network interruptions might occur during migration depending on the physical network environment. Proceed with caution for service-sensitive VMs.
- If the selected VM has snapshots, storage migration while the VM is running deletes these snapshots. Proceed with caution. To keep snapshots, shut down the VM before migration.

What's next

After migration, you can re-attach the detached data disks, ISOs, or peripheral devices to the VM.

5.7.2.4 Migrate Data Disks

Prerequisites

- Make sure that the source and destination data storage are attached to the same cluster.
- SAN data disks can be migrated while the VM is running or shut down. NFS and nSDS data disks can be migrated only after the VM is shut down.

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. Select the target VM and click **Actions** > **Migration Management** > **Change Data Storage**.
3. In the risk confirmation dialog, review the migration risks, select **I understand the risks above**, and click **OK**.
4. In the Migration Scope section of the **Change Data Storage** dialog, select **Data Disk**.
5. Select the data disks to migrate and configure a destination for each disk.

Option	Description
Migrate All	Selects all data disks that can be migrated.
Cancel All	Clears the selection of all data disks.
Migrate	Specifies whether to migrate the corresponding data disk.
Destination Storage	Specifies the destination data storage for the selected data disk. You can specify a different destination for each data disk.
Storage Pool	Specifies the destination storage pool when the destination is nSDS distributed storage.
Migration Encryption	Encrypts data transferred during migration. If VM encryption is enabled, this option is enabled automatically and cannot be disabled. If VM encryption is not enabled, this option is disabled by default and can be enabled as needed. The disk encryption state does not affect this option. For more information, see *Interoperability with Other Features*.

If the VM has many data disks, search for a disk by name.

6. Review the configuration and click **OK**.

5.7.3 Change Host and Data Storage

5.7.3.1 Hot Migrate a VM Across Data Storage of the Same Type

Prerequisites

- Make sure that the data storage where the VM resides meets the storage combination requirements.

  Supported storage combinations:

  - Hot migration from local storage to local storage.
  - Hot migration from nSDS distributed storage to nSDS distributed storage.
  - Hot migration from NFS storage to NFS storage.
  - Hot migration from SAN storage to SAN storage.
- If the VM has shared disks, ISOs, or peripheral devices attached, detach them before migration .

- If vNUMA is enabled for the VM, make sure that the pNUMA topology of the destination host is consistent with that of the source host.
- If the VM has VF NICs attached, make sure that the following requirements are met before hot migration:
  - VMTools is installed on the VM and is running properly.
  - If the driver version inside the VM is too old, migration might fail. Check and update the driver version as needed.
- If the selected VM is associated with VM scheduling policies, modify the policies based on your service requirements to avoid policy conflicts.
- For hot migration across data storages of the same type, only full-VM migration is supported.
- For hot migration from one nSDS distributed storage to another nSDS distributed storage, make sure that the monitoring nodes of the two distributed storage systems can communicate with each other.

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. Select the target VM and click **Actions** > **Migration Management** > **Change Host and Data Storage**.
3. In the **Change Host and Data Storage** dialog, choose the destination data storage and destination host, and configure the migration bandwidth and **Migration Encryption** as needed.
   - **Migration Encryption** encrypts memory, runtime state, and disk data transferred during hot migration. If VM encryption is enabled, this option is enabled automatically and cannot be disabled. If VM encryption is not enabled, this option is disabled by default and can be enabled as needed. The disk encryption state does not affect this option. For more information, see *Interoperability with Other Features*.
   - For hot migration from one nSDS distributed storage to another nSDS distributed storage, you can specify the storage pool for disk 1 and choose whether to migrate data disks.
   - If the VM has been running under a high service load for a long time and the service can tolerate a moderate performance impact, you can enable auto-converge to improve the migration success rate.
4. Review the configuration and click **OK**.



Note:

- If the selected VM has VF NICs attached, short network interruptions might occur during migration depending on the physical network environment. Proceed with caution for service-sensitive VMs.
- If the selected VM has snapshots, storage migration while the VM is running deletes these snapshots. Proceed with caution. To keep snapshots, shut down the VM before migration.

5.7.3.2 Cold Migrate a VM Across Data Storage of the Same Type

Prerequisites

- Make sure that the data storage where the VM resides meets the storage combination requirements.

  Supported storage combinations:

  - Cold migration from nSDS distributed storage to nSDS distributed storage.
  - Cold migration from NFS storage to NFS storage.
- If the VM has data disks, shared disks, ISOs, or peripheral devices attached, detach them before migration.
- If the selected VM is associated with VM scheduling policies, modify the policies based on your service requirements to avoid policy conflicts.
- For cold migration from one nSDS distributed storage to another nSDS distributed storage, if a cold migration network is configured for the destination distributed storage, the VM is migrated through the cold migration network. Otherwise, the management network is used.
- For cold migration from one NFS storage to another NFS storage, the destination NFS storage must be attachable to the cluster where the VM to be migrated resides.

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. Select the target VM and click **Shut Down**.
3. After the VM is shut down, click **Actions** > **Migration Management** > **Change Host and Data Storage**.
4. In the **Change Host and Data Storage** dialog, choose the destination data storage and destination host, and configure **Migration Encryption** as needed.

   Migration Encryption encrypts data transferred during migration. If VM encryption is enabled, this option is enabled automatically and cannot be disabled. If VM encryption is not enabled,

this option is disabled by default and can be enabled as needed. The disk encryption state does not affect this option. For more information, see *Interoperability with Other Features*.

5. Review the configuration and click **OK**.

5.7.3.3 Cold Migrate a VM Across Different Data Storage Types

Prerequisites

- Make sure that the data storage where the VM resides meets the storage combination requirements.

  Supported storage combinations:

  - Cold migration between nSDS distributed storage and SAN storage.
  - Cold migration between nSDS distributed storage and NFS storage.
  - Cold migration between local storage and SAN storage.
  - Cold migration between local storage and nSDS distributed storage.
  - Cold migration between local storage and NFS storage.
  - Cold migration between SAN storage and NFS storage.
- If the VM has shared disks, ISOs, or peripheral devices attached, detach them before migration .
- If vNUMA is enabled for the VM, make sure that the pNUMA topology of the destination host is consistent with that of the source host.
- If the VM has VF NICs attached, make sure that the following requirements are met before migration:
  - VMTools is installed on the VM and is running properly.
  - If the driver version inside the VM is too old, migration might fail. Check and update the driver version as needed.
- If the selected VM is associated with VM scheduling policies, modify the policies based on your service requirements to avoid policy conflicts.
- For hot migration across data storages of different types, only full-VM migration is supported.

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. Select the target VM and click **Actions** > **Migration Management** > **Change Host and Data Storage**.

3. In the **Change Host and Data Storage** dialog, choose the destination data storage and destination host, and configure the migration bandwidth and **Migration Encryption** as needed.
   - **Migration Encryption** encrypts data transferred during migration. If VM encryption is enabled, this option is enabled automatically and cannot be disabled. If VM encryption is not enabled, this option is disabled by default and can be enabled as needed. The disk encryption state does not affect this option. For more information, see *Interoperability with Other Features*.
   - For migration from SAN storage, local storage, or NFS storage to nSDS distributed storage, you can specify the storage pool for disk 1.
   - If the VM has been running under a high service load for a long time and the service can tolerate a moderate performance impact, you can enable auto-converge to improve the migration success rate.
4. Review the configuration and click **OK**.

Note:

- If the selected VM has VF NICs attached, short network interruptions might occur during migration depending on the physical network environment. Proceed with caution for service-sensitive VMs.
- If the selected VM has snapshots, storage migration while the VM is running deletes these snapshots. Proceed with caution. To keep snapshots, shut down the VM before migration.

5.7.3.4 Hot Migrate a VM Across Storage Pools in the Same nSDS Distributed Storage

Prerequisites

- Make sure that the VM resides on nSDS distributed storage.
- If the VM has shared disks, ISOs, or peripheral devices attached, detach them before migration.
- If vNUMA is enabled for the VM, make sure that the pNUMA topology of the destination host is consistent with that of the source host.
- If the VM has VF NICs attached, make sure that the following requirements are met before hot migration:
  - VMTools is installed on the VM and is running properly.

- If the driver version inside the VM is too old, migration might fail. Check and update the driver version as needed.

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. Select the target VM and click **Actions** > **Migration Management** > **Change Host and Data Storage**.
3. In the **Change Host and Data Storage** dialog, choose the destination data storage and destination host, and configure the migration bandwidth and **Migration Encryption** as needed.
   - **Migration Encryption** encrypts memory, runtime state, and disk data transferred during hot migration. If VM encryption is enabled, this option is enabled automatically and cannot be disabled. If VM encryption is not enabled, this option is disabled by default and can be enabled as needed. The disk encryption state does not affect this option. For more information, see *Interoperability with Other Features*.
   - For hot migration across storage pools in the same nSDS distributed storage, you can specify the storage pool for disk 1 and choose whether to migrate data disks.
   - If the VM has been running under a high service load for a long time and the service can tolerate a moderate performance impact, you can enable auto-converge to improve the migration success rate.
4. Review the configuration and click **OK**.

Note:

- If the selected VM has VF NICs attached, short network interruptions might occur during migration depending on the physical network environment. Proceed with caution for service-sensitive VMs.
- If the selected VM has snapshots, storage migration while the VM is running deletes these snapshots. Proceed with caution. To keep snapshots, shut down the VM before migration.

5.7.3.5 Cold Migrate a VM Across Storage Pools in the Same nSDS Distributed Storage

Prerequisites

If the VM has data disks, shared disks, ISOs, or peripheral devices attached, detach them before migration.

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. Select the target VM and click **Shut Down**.
3. After the VM is shut down, click **Actions** > **Migration Management** > **Change Host and Data Storage**.
4. In the **Change Host and Data Storage** dialog, choose the destination data storage and destination host, and configure **Migration Encryption** as needed.

   Migration Encryption encrypts data transferred during migration. If VM encryption is enabled, this option is enabled automatically and cannot be disabled. If VM encryption is not enabled, this option is disabled by default and can be enabled as needed. The disk encryption state does not affect this option. For more information, see *Interoperability with Other Features*.
5. Review the configuration and click **OK**.

5.8 Manage VMs with Snapshots

VM snapshots preserve the state of a VM at a specific point in time. They are suitable for keeping a temporary rollback point before system changes, application upgrades, or configuration adjustments. You can create snapshots manually or use snapshot policies to create snapshots on a schedule.

A snapshot is not a backup and does not replace a disaster recovery solution. In production environments, use snapshots carefully based on service load, storage capacity, and retention requirements.

A snapshot inherits the encryption state of its disks. Snapshot data created from an encrypted disk remains encrypted and does not require separate snapshot encryption settings.

Snapshot Types

- **Manual snapshot**: Manually create a snapshot for the entire VM at any time. For more information, see *Create a Snapshot*.
- **Scheduled snapshot**: Use a snapshot policy to automatically create snapshots at specified times. For more information, see *Create a Snapshot Policy*.
- **Automatic snapshot**: The system triggers a one-time automatic snapshot in specific scenarios, such as system reset and VM cloning. When an image is created, the system automatically creates a temporary snapshot and deletes it after image creation is complete.

Scenarios

- Fast fault recovery: If an exception occurs in a production environment, you can use snapshot rollback to quickly restore the VM to a normal state. Snapshot rollback is intended for temporary recovery from unexpected faults.
- Data development: Create snapshots for production data to provide near-real production data for data mining, report query, development, and testing.
- Improved operation fault tolerance: Before major operations such as system upgrades or service data migration, create one or more snapshots. If an error occurs during the operation, you can restore the system data to a normal state from a snapshot.

5.8.1 Create a Snapshot

You can create one or more snapshots for a VM to retain its temporary state at a specific point in time, allowing quick rollback if a fault occurs.

Prerequisites

- If the VM has shared disks attached, snapshots are not supported.
- Memory snapshots are not supported for VMs that use distributed storage.
- Memory snapshots are not supported for VMs in Compatible CPU mode.
- Before you create a memory snapshot, make sure that the VM is running and that attached peripheral devices are detached.
- Too many VM snapshots can affect VM performance, data security, and data storage capacity. For long-term data protection, use dedicated backup or protection features.

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. Select the target VM and click **Actions** > **Snapshot and Image** > **Create Snapshot**.
3. In the **Create Snapshot** dialog, configure the following options:

Option	Description
Name	Enter a name for the VM snapshot.
Description	Enter a brief description for the snapshot.
Memory Snapshot	A memory snapshot preserves the real-time state of the VM.

4. Review the configuration and click **OK**.

Note:
Snapshot data created from an encrypted disk remains encrypted. No additional encryption settings are required.

Note:
To ensure memory consistency, the VM is paused briefly during memory snapshot creation.

5.8.2 Revert to a Snapshot

Reverting a VM to the selected snapshot clears the current data and restores the VM data to the state captured when the selected snapshot was created.

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. Select the target VM and click **Snapshot**.
3. On the **Snapshot** tab, select the snapshot point to restore and click **Revert**.
4. In the **Revert Snapshot?** dialog, configure the following options:

Option	Description
Recovery Method	Supports full recovery and custom recovery. Reverting to a snapshot rolls back the current data. • Full Recovery: Restores the data and disk sequence of the VM system disk and data disks to the snapshot point. Disks detached after snapshot creation might be attached again, and disks attached after snapshot creation might be detached. • Custom Recovery: Select disks for recovery. Only the data of the selected disks is restored, and the disk sequence is not restored. **Note:** If a disk in the snapshot has been detached, recovery of that disk is not currently supported.
Power Status	Choose whether to automatically power on the VM after snapshot revert.

5. Review the configuration and click **OK**.

Note:

- During snapshot revert, the system automatically powers off the VM, clears the current data , and rolls back VM data to the selected snapshot point.
- If the number of VM NICs or CD/DVD drives differs from the configuration recorded in the snapshot, NICs and CD/DVD drives removed after snapshot creation are attached again, and newly attached NICs and CD/DVD drives are removed.
- When reverting to a memory snapshot, if the IP address or MAC address in the snapshot conflicts with an existing VM, resolve the conflict before the revert.

  If you force the revert, the platform handles conflicts as follows:

  - IP address conflict: The MAC address in the snapshot continues to be used. You must manually adjust the IP address after the revert.
  - MAC address conflict: The MAC address in the snapshot continues to be used, but the related NIC is disabled after the revert.

5.8.3 Create a VM from a Snapshot

Create a new VM directly from a snapshot. This method does not overwrite the existing VM data.

Prerequisites

- Make sure that the VM has at least one successfully created snapshot point.
- Make sure that the platform has sufficient compute, storage, and network resources for the new VM.
- If the selected snapshot is a memory snapshot, creating a VM from the snapshot is not supported.
- If the snapshot contains encrypted disks, make sure that the key provider and keys protecting the snapshot are available. To enable VM encryption or encrypt an unencrypted disk, make sure that the default key provider is available.

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. Select the target VM and click **Snapshot**.
3. On the **Snapshot** tab, select the snapshot point and click **New Virtual Machine**.
4. In the **New Virtual Machine from Snapshot** dialog, configure the following options:

Snapshot Information

Option	Description
Snapshot	Displays the selected snapshot point.

Basic Information

Option	Description
Name	VM name.
Quantity	The default value is 1 and cannot be changed.
Location	Host or cluster where the VM resides.
OS	VM operating system, including Linux and Windows.
Power State	Choose whether to automatically power on the VM after it is created.
VM Encryption	Follows the current VM encryption state of the source VM. If VM encryption is currently enabled for the source VM, this option remains enabled and cannot be changed. Otherwise, the option is disabled by default and can be enabled as needed.

Hardware Information

Option	Description
CPU	Supports adjusting the total number of CPU cores.
Memory	Supports adjusting the memory size.
Disk	Displays the disk configuration recorded in the snapshot. Except for disk encryption, the configuration cannot be changed. Disk encryption is handled as follows: • An encrypted disk in the snapshot remains encrypted, and disk encryption cannot be disabled. • An unencrypted disk in the snapshot remains unencrypted by default. You can enable disk encryption as needed.

Option	Description
	• If you manually enable VM encryption, disk encryption is also enabled for unencrypted disks.
NIC	Supports adjusting the port group, MAC address, IP address, DNS assignment mode, and security group. Click **Add NIC** to add a NIC to the VM. You can customize the network address and features of the new NIC.

5. Review the configuration and click **OK**.

What's next

Some VM configurations depend on VMTools. After the VM is created, install VMTools so that related configurations can take effect. For more information about VMTools, see *Install and Use VMTools*.

5.8.4 View Snapshots

The platform provides a centralized page for managing all snapshots. You can also view snapshots of a specific VM on the VM's **Snapshot** tab.

Procedure

1. In the navigation pane, choose **Data Protection** > **Snapshot**.
2. The **Snapshot** page displays snapshot information in a tree hierarchy and supports sorting by snapshot quantity or snapshot capacity.
   - After you select a VM, the details area displays all snapshots of the VM, including snapshot name, capacity, creation time, and available actions.
   - After you select a specific snapshot of a VM, the details area displays the available actions, basic information, and snapshot information of the snapshot.

5.8.5 Delete a Snapshot

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. Select the target VM and click **Snapshot**.
3. On the **Snapshot** tab, select the snapshot to delete and click **Actions** > **Delete**.
4. Confirm the selected snapshot information and risk notice, and then click **OK**.

Note:

- Deleting the current snapshot also deletes snapshots on its child branches.
- Deleting a snapshot consumes I/O. Do not restart the management node service during deletion.
- Snapshot data is merged before the snapshot is deleted. This process requires a certain amount of data storage capacity. If data storage capacity is insufficient, snapshot deletion might fail.

5.8.6 Snapshot Policy Management

5.8.6.1 Create a Snapshot Policy

Associate a snapshot policy with a specified VM to create snapshots automatically on a schedule.

Prerequisites

- A VM can be associated with only one snapshot policy. A VM that already has a snapshot policy cannot be associated with another snapshot policy.
- If a VM has shared disks or RDM disks attached, it cannot be associated with a snapshot policy .
- Make sure that the VM resides on nSDS distributed storage.

Procedure

1. In the navigation pane, choose **Data Protection** > **Snapshot Policy**.
2. On the **Snapshot Policy** page, click **New Snapshot Policy**.
3. In the **New Snapshot Policy** dialog, configure the following options:

Option	Description
Name	Snapshot policy name.
Description	Enter a brief description for the snapshot policy.
Scheduled Snapshot Cycle	Set the snapshot creation frequency. Snapshots can be created weekly or monthly, and creation time can be configured down to the minute.
Start Time	Set the time when the snapshot policy starts to run.

Option	Description
End Time	Set the end time of the snapshot policy. The value can be Never or Custom.
Retained Snapshots	Set the maximum number of snapshots to retain. Snapshots that exceed the limit are automatically deleted.
Associated VM	After VMs are associated, the system automatically creates snapshots for the associated VMs at the specified time when the snapshot policy takes effect.

4. Review the configuration and click **OK**.

5.8.6.2 Enable or Disable a Snapshot Policy

Manage the snapshot policy state by enabling or disabling the policy. After a snapshot policy is disabled, scheduled automatic snapshots are paused until the policy is enabled again.

Procedure

1. In the navigation pane, choose **Data Protection** > **Snapshot Policy**.
2. On the **Snapshot Policy** page, select the target snapshot policy and click **Actions** > **Enable/ Disable**.

5.8.6.3 Modify a Snapshot Policy

You can modify a snapshot policy, including its scheduled snapshot cycle, start time, end time, retained snapshot count, and associated VMs.

Procedure

1. In the navigation pane, choose **Data Protection** > **Snapshot Policy**.
2. On the **Snapshot Policy** page, select the target snapshot policy and click **Actions** > **Modify Configuration**.
3. In the **Modify Configuration** dialog, modify the snapshot policy as needed.
4. Review the updated configuration and click **OK**.

5.8.6.4 Delete a Snapshot Policy

Procedure

1. In the navigation pane, choose **Data Protection** > **Snapshot Policy**.
2. On the **Snapshot Policy** page, select the target snapshot policy and click **Actions** > **Delete**.

5.8.7 Snapshot Best Practices

When you use snapshots in a production environment, follow these recommendations:

1. In a production environment, keep the number of snapshots for a single disk within five whenever possible. Too many snapshots can affect VM or disk I/O performance, data security, and data storage capacity.
2. In a production environment, avoid creating snapshots for high-I/O VMs to ensure data integrity . When high I/O operations are running inside a VM, some data in memory might not yet have been written to disk and therefore cannot be saved in the snapshot.
3. Snapshot capacity usage:
   - In local storage or centralized storage scenarios, a new incremental snapshot uses only a small amount of storage space, while a new full snapshot uses storage space equivalent to one full copy.
   - In distributed storage scenarios, creating a snapshot does not consume additional storage space by itself. However, after a snapshot is created, write operations to the original disk might trigger Copy-On-Write, causing each snapshot to consume storage space equivalent to the original disk.
4. Before snapshot rollback, create a snapshot for the disk to protect the current disk data state.
5. Impact of creating snapshots on current services:

   In most cases, creating a snapshot does not affect current services. However, creating a full snapshot generates additional network I/O on the storage network. In local storage or centralized storage scenarios, disk I/O bandwidth usage is more noticeable, while disk IOPS are generally not affected.
6. Impact of deleting snapshots on current services:

   In most cases, deleting a snapshot also deletes snapshots on its child branches and merges the data into the disk. This process generates additional disk I/O bandwidth, and service I/O might slow down slightly.

5.9 Maintenance and Troubleshooting

VM maintenance and troubleshooting help administrators handle exceptions that occur while VMs are running. You can configure a VM failure handling policy so that the platform performs a predefined action when a VM exception is detected.

Before you configure a failure handling policy, make sure that the VM meets the runtime requirements for the policy to take effect, and evaluate the impact of the automatic action on services.

5.9.1 Configure a VM Failure Response Policy

A VM failure refers to a Windows VM blue screen or a Linux VM becoming unresponsive during operation due to excessive load or other issues. You can configure an automatic response policy for VM failures.

Prerequisites

- Failure response policies are supported only for VMs that use the x86_64 CPU architecture.
- Before you configure a failure response policy, install VMTools on the VM and make sure that VMTools is running properly. For more information about VMTools, see *Install and Use VMTools*.
- The failure response policy takes effect immediately after configuration and does not require a VM restart.

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. Select the target VM and click **Actions** > **Advanced Settings** > **Modify VMTools**.
3. In the **Modify VMTools Settings** dialog, modify the VM failure response policy as needed.
   - No Action (default): Keep the current state after the failure and take no action.
   - Reboot: Automatically reboot the VM after the failure. By default, after five reboot attempts within 30 minutes, the system stops rebooting the VM and keeps it in the failed state.
   - Shut Down: Automatically shut down the VM after the failure.

   Note:
   For Linux VMs, after this feature is enabled, the `kdump` module is disabled at the next startup.
4. Click **OK**.

6 Storage

6.1 Overview

NexaSphere Storage manages the data storage resources that provide read and write capacity for virtual machines. You can add local storage, NFS storage, SAN storage, or distributed storage based on the deployment requirements, and then maintain the state, status, configuration, capacity, and deletion operations of the data storage resources.

You can also use this book to deploy a distributed storage service quickly or take over an existing distributed storage service on NexaSphere.

6.2 Data Storage Types and Resource Relationships

Data storage provides storage space for virtual machines in NexaSphere. After data storage is attached to a cluster, virtual machines in the cluster can create and use resources such as system disks, data disks, and snapshots on the corresponding data storage.

NexaSphere supports local storage, NFS storage, SAN storage, and distributed storage. Each type has different backend resources, attachment relationships, and suitable scenarios. SAN storage is created from LUNs provided by block storage targets. Distributed storage is created from available distributed storage resources or storage pools.

Image storage provides image services for clusters in a data center. Different image storage types must be used with compatible data storage types. When planning data storage, consider both the image storage type and the cluster attachment relationship.

6.2.1 Data Storage Types

Local Storage

Local storage uses the local disks of hosts as data storage resources. It is suitable when you want to make full use of local storage capacity on hosts.

NFS Storage

NFS storage is accessed through a Network File System. You can use custom mount parameters to control the NFS mount behavior.

SAN Storage

SAN storage uses shared block storage and supports thin provisioning and thick provisioning as storage allocation strategies. SAN storage can be created from LUNs provided by iSCSI storage, FC storage, or NVMe storage.

Distributed Storage

Distributed storage uses distributed block storage. nSDS, ZHPS, and ZBS distributed storage can be added to NexaSphere. After it is added, NexaSphere uses the distributed storage as data storage to provide storage capacity for clusters and virtual machines.

6.2.2 Storage Devices and SAN Storage Targets

SAN storage targets provide the backend resources for SAN data storage. In NexaSphere, you can add or synchronize iSCSI, FC, and NVMe storage targets, and then use the LUNs discovered from those targets to add SAN data storage.

Before you add SAN data storage, add or synchronize the corresponding SAN storage target and make sure that the LUNs are provisioned, mapped, and reachable on the storage side. When SAN data storage is added, NexaSphere initializes or registers the LUNs discovered from the storage target as SAN data storage.

SAN Storage Target Types

- iSCSI storage: Accesses block storage devices over IP networks.
- FC storage: Accesses block storage devices over Fibre Channel networks.
- NVMe storage: Accesses block storage devices over the NVMe-oF protocol.

Usage Notes

- A LUN that is not attached to any virtual machine can be added as SAN storage.
- After a LUN is added as SAN storage, the LUN cannot be used for other purposes.
- A LUN that is not added as SAN storage can be attached to virtual machines. One LUN can be attached to multiple virtual machines, and one virtual machine can attach multiple LUNs.

6.2.3 Cluster, Data Storage, and Image Storage Relationships

Clusters and Data Storage

Clusters and data storage are both sub-resources in a data center. Data storage must be attached to a cluster before it can provide storage services to virtual machines in the cluster. The number of data storage resources that can be attached to a cluster varies by storage type.

Resource A	Resource B	Quantitative Relationship Between Resource A and Resource B
Cluster	Local Storage	1:n
	NFS Storage	1:n
	SAN Storage	1:n
	nSDS Distributed Storage	1:1
	ZHPS Distributed Storage	1:1
	ZBS Distributed Storage	1:1
	Local Storage + NFS Storage	1:(1 Local Storage + 1 NFS Storage)
	Local Storage + SAN Storage	1:(n Local Storage + n SAN Storage)
	nSDS Distributed Storage + Local Storage	1:(1 nSDS Distributed Storage + 3 Local Storage)
	nSDS Distributed Storage + SAN Storage	1:(1 nSDS Distributed Storage + n SAN Storage)
	NFS Storage + SAN Storage	1:(n NFS Storage + n SAN Storage)

Data Storage and Image Storage

After image storage is attached to a data center, it provides image services for clusters in the data center. Different image storage types must be used with compatible data storage types to support operations such as creating virtual machines in a cluster.

Image Storage Type	Data Storage Type	Supports Creating VMs
Standalone Image Storage	Local Storage	Yes
	NFS Storage	Yes
	nSDS Distributed Storage	Yes
	SAN Storage	Yes
	ZHPS Distributed Storage	Yes
	ZBS Distributed Storage	Yes
Distributed Image Storage	Local Storage	No
	NFS Storage	No

Image Storage Type	Data Storage Type	Supports Creating VMs
	nSDS Distributed Storage	Yes
	SAN Storage	No

6.3 SAN Storage Target Management

6.3.1 Add an iSCSI Storage Target

Prerequisites

You have allocated and mapped the LUNs on the storage side.

Procedure

1. In the navigation pane, choose **Inventory** > **Data Storage**.
2. Select the target data center and click **Storage Target**.
3. On the **Storage Target** tab, select **iSCSI Storage**.
4. Click **Add iSCSI Storage**.
5. In the **Add iSCSI Storage** dialog, configure the following options:

Option	Description
Name	Enter a name for the iSCSI storage.
IP Address	Enter the IP address of the iSCSI storage server.
Port	Enter the target port for the iSCSI storage. Default: 3260.
Cluster	Select the cluster where the iSCSI storage will be attached. You can also attach the iSCSI storage to a cluster after adding it.  **Note:** If you want to use LUNs provided by the iSCSI storage, attach the iSCSI storage to the cluster where the SAN storage is attached.
CHAP Username	Enter the CHAP authentication username.
CHAP Password	Enter the CHAP authentication password.

6. Review the configuration and click **OK**.

What's next

You can add iSCSI LUNs as SAN storage or pass them through to virtual machines. LUNs that have been added as SAN storage cannot be passed through to virtual machines.

- To add iSCSI LUNs as SAN storage, see *Add SAN Storage*.
- For more information about passing through LUNs to virtual machines by using RDM disks, see *Create a Virtual Machine*.

6.3.2 Synchronize an FC Storage Target

Prerequisites

- You have set up the FC storage in advance.
- Make sure network connectivity between the host and storage is established.

Procedure

1. In the navigation pane, choose **Inventory** > **Data Storage**.
2. Select the target data center and click **Storage Target**.
3. On the **Storage Target** tab, select **FC Storage**.
4. Click **Sync Device Info**.

What's next

You can add FC LUNs as SAN storage or pass them through to virtual machines. LUNs that have been added as SAN storage cannot be passed through to virtual machines.

- To add FC LUNs as SAN storage, see *Add SAN Storage*.
- For more information about passing through LUNs to virtual machines by using RDM disks, see *Create a Virtual Machine*.

6.3.3 Add an NVMe Storage Target

Prerequisites

- You have set up the NVMe storage in advance.
- Make sure network connectivity between the host and storage is established.

Procedure

1. In the navigation pane, choose **Inventory** > **Data Storage**.
2. Select the target data center and click **Storage Target**.
3. On the **Storage Target** tab, select **NVMe Storage**.
4. Click **Add NVMe Storage**.

5. In the **Add NVMe Storage** dialog, configure the following options:

Option	Description
Name	Enter a name for the NVMe storage.
Transmission Mode	Supports RDMA and TCP transmission modes.
IP Address	IP address of the NVMe storage server.
Port	Port of the NVMe storage. Default: 4420.
Cluster	The cluster where the NVMe storage will be attached.

6. Review the configuration and click **OK**.

What's next

You can add NVMe LUNs as SAN storage or pass them through to virtual machines. LUNs that have been added as SAN storage cannot be passed through to virtual machines.

- To add NVMe LUNs as SAN storage, see *Add SAN Storage*.
- For more information about passing through LUNs to virtual machines by using RDM disks, see *Create a Virtual Machine*.

6.4 Add Data Storage

6.4.1 Add Local Storage

Prerequisites

- Plan the mount path for the local storage on the host disk in advance.
- Check the quantity and type limitations between the cluster and the data storage to be attached. For more information, see *Clusters and Data Storage*.

Procedure

1. In the navigation pane, choose **Inventory** > **Data Storage**.
2. Select the target data center and click **Actions** > **Add Data Storage**.
3. In the **Add Data Storage** dialog, configure the following options:

   Basic Information

Option	Description
Name	Enter a name for the local storage.

Option	Description
Description	Enter a brief description for the data storage.
Type	Select **Local Storage**.
Data Center	Displays the data center that contains the local storage.

Configurations

Option	Description
Cluster	Select the cluster where the local storage will be attached.
Mount Path	Enter the mount path for the local storage on the host disk.  **Note:** System directories such as `/`, `/dev`, `/proc`, `/sys`, `/usr/bin`, and `/bin` cannot be used. Using system directories might cause hosts to work improperly.
Host Disk	Select the unmounted or unpartitioned free disks on hosts that are in a "Connected" status under the selected cluster.  **Note:** Configuring a free disk formats the selected disk and completely clears all partitions, file systems, and data on the disk. If host disks are left unselected, the specified mount path will be added as the local directory.

4. Review the configuration and click **OK**.

6.4.2 Add NFS Storage

Prerequisites

- Plan the mount path of the NFS storage in advance and configure access permissions for the corresponding directory on the NFS Server.
- Check whether the NFS Server supports mount parameters. If it does, specify the parameters when adding NFS storage to optimize the performance, security, and reliability of the network file system.

- Plan a dedicated storage network in advance to avoid network congestion.
- Check the quantity and type limitations between the cluster and the data storage to be attached. For more information, see *Clusters and Data Storage*.

Procedure

1. In the navigation pane, choose **Inventory** > **Data Storage**.
2. Select the target data center and click **Actions** > **Add Data Storage**.
3. In the **Add Data Storage** dialog, configure the following options:

   Basic Information

Option	Description
Name	Enter a name for the NFS storage.
Description	Enter a brief description for the data storage.
Type	Select **NFS**.
Data Center	Displays the data center that contains the NFS storage.

   Configurations

Option	Description
Cluster	Select the cluster where the NFS storage will be attached.
Mount Path	Enter the shared directory URL on the NFS Server. Supported format: *NFS_Server_IP:/NFS_Share_folder*.  **Note:** System directories such as `/`, `/dev`, `/proc`, `/sys`, `/usr/bin`, and `/bin` cannot be used. Using system directories might cause hosts to work improperly.
Mount Option	Specify mount parameters on the NFS Server. For parameter values, see the `-o` option of `mount`. If the parameters configured here conflict with the settings on the NFS Server, the NFS Server settings take precedence.

Option	Description
Storage Network	Enter the CIDR of the storage network specified for the NFS storage.

4. Review the configuration and click **OK**.

6.4.3 Add SAN Storage

Prerequisites

- LUNs used by SAN storage are provided by iSCSI storage, FC storage, or NVMe storage. For more information, see *Add an iSCSI Storage Target*, *Synchronize an FC Storage Target*, and *Add an NVMe Storage Target*.
- Check the quantity and type limitations between the cluster and the data storage to be attached. For more information, see *Clusters and Data Storage*.

Procedure

1. In the navigation pane, choose **Inventory** > **Data Storage**.
2. Select the target data center and click **Actions** > **Add Data Storage**.
3. In the **Add Data Storage** dialog, configure the following options:

   Basic Information

Option	Description
Name	Enter a name for the SAN storage.
Description	Enter a brief description for the data storage.
Type	Select **SAN Storage**.
Data Center	Displays the data center that contains the SAN storage.

   Configurations

Option	Description
Cluster	Select the cluster where the SAN storage will be attached.
Default Provisioning Method of Storage Space	Specify the allocation method for disk storage space. Default: Thick Provision.

Option	Description
	• **Thick Provision**: Pre-allocates the required storage space to provide sufficient capacity and ensure storage performance. • **Thin Provision**: Allocates storage space according to actual usage for higher storage utilization.
Storage Addition Method	Select a method to add storage. • **New**: Creates new data storage by initializing unused LUNs into data storage that can be used by the platform. You can choose whether to clear existing data on the LUNs when adding them. • **Register**: Imports LUNs that already contain data, such as storage migrated from another platform or replicated through disaster recovery, into the platform. The platform identifies the SAN storage structure and repairs metadata. During registration, all LUNs under the SAN storage are added automatically.
LUN	Select LUNs provided by iSCSI storage, FC storage, or NVMe storage.
Storage Network	Specify a storage network for shared storage. The platform uses this storage network to check the health status of virtual machines.
Cleanse LUN	Choose whether to forcibly clear existing data on the LUNs, such as file system, RAID, or partition table labels. By default, the LUNs are not cleansed.  **Note:** If data or partitions exist on the LUNs, adding LUNs or attaching data storage might fail.
Storage UUID	When you select **Register** as the storage addition method, specify how to handle the storage UUID. • **Keep**: Retains the original UUID of the SAN storage.

Option	Description
	• **Reset**: Regenerates a UUID for the SAN storage. This option applies to scenarios such as UUID conflicts after cloning or snapshots, or disaster recovery drills where the storage must be decoupled from the production environment. After reset, the original UUID becomes invalid. If other hosts still use the original UUID to access the storage, I/O errors occur.

4. Review the configuration and click **OK**.

6.4.4 Add nSDS Distributed Storage

Prerequisites

- You have set up the nSDS distributed storage on the storage side. To use specified image cache pools and data disk pools, create the corresponding storage pools in the distributed storage cluster in advance.
- Plan a dedicated storage network in advance to avoid network congestion.
- Check the quantity and type limitations between the cluster and the data storage to be attached. For more information, see *Clusters and Data Storage*.

Procedure

1. In the navigation pane, choose **Inventory** > **Data Storage**.
2. Select the target data center and click **Actions** > **Add Data Storage**.
3. In the **Add Data Storage** dialog, configure the following options:

   Basic Information

Option	Description
Name	Enter a name for the distributed storage.
Description	Enter a brief description for the data storage.
Type	Select **Distributed Storage nSDS**.
Data Center	Displays the data center that contains the nSDS distributed storage.

   Configurations

Option	Description
Cluster	Select the cluster where the nSDS distributed storage will be attached.
Key Authentication	Choose whether to use key authentication for distributed storage. Default: Enabled. **Note:** Make sure the key authentication setting on the distributed storage side is consistent with this option. If key authentication is not disabled on distributed storage but is disabled here, VM creation might fail, and vice versa.
Monitoring Node	Add monitoring nodes and configure the management IP address, SSH port, username, and password of each monitoring node.
Image Cache Pool	Specify a storage pool for image caches. If left empty, the system creates one automatically.
Storage Pool	Specify a storage pool for data volumes. If left empty, the system creates one automatically.
Storage Network	Enter the CIDR of the storage network specified for the nSDS distributed storage.

4. Review the configuration and click **OK**.

6.4.5 Add ZHPS Distributed Storage

ZHPS Distributed Storage connects to high-performance SSD distributed storage in vhost-user mode.

Prerequisites

- You have set up the ZHPS distributed storage on the storage side.
- Check the quantity and type limitations between the cluster and the data storage to be attached. For more information, see *Clusters and Data Storage*.
- Make sure the platform has a valid Advanced Edition license.

Procedure

1. In the navigation pane, choose **Inventory** > **Data Storage**.

2. Select the target data center and click **Actions** > **Add Data Storage**.
3. In the **Add Data Storage** dialog, configure the following options:

   Basic Information

Option	Description
Name	Enter a Name for the ZHPS distributed storage.
Description	Enter a brief description for the data storage.
Type	Select **Distributed Storage ZHPS**.
Data Center	Displays the data center that contains the ZHPS distributed storage.

   Configurations

Option	Description
Cluster	Select the cluster where the ZHPS distributed storage will be attached.
IP Address	Enter the management address of the ZHPS distributed storage.
Port	Enter the port of the ZHPS distributed storage management address. Default: 443.
Username	Enter the username for the ZHPS distributed storage management address.
Password	Enter the password of the username.
Storage Pool	After the connection test succeeds, specify a storage pool that has been created on the storage side.

4. Review the configuration and click **OK**.

6.4.6 Add ZBS Distributed Storage

ZBS Distributed Storage connects to high-performance distributed block storage through the CBD interface.

Prerequisites

- You have set up the ZBS distributed storage on the storage side.

- Check the quantity and type limitations between the cluster and the data storage to be attached. For more information, see *Clusters and Data Storage*.
- Make sure the platform has a valid Advanced Edition license.

Procedure

1. In the navigation pane, choose **Inventory** > **Data Storage**.
2. Select the target data center and click **Actions** > **Add Data Storage**.
3. In the **Add Data Storage** dialog, configure the following options:

   Basic Information

Option	Description
Name	Enter a Name for the ZBS distributed storage.
Description	Enter a brief description for the data storage.
Type	Select **Distributed Storage ZBS**.
Data Center	Displays the data center that contains the ZBS distributed storage.

   Configurations

Option	Description
Cluster	Select the cluster where the ZBS distributed storage will be attached.
MDS Node	Add MDS nodes and specify the management IP address, SSH port, username, and password of each MDS node.
Storage Pool	Enter the pool name after creating the storage pool on the storage side.

4. Review the configuration and click **OK**.

6.5 State, Status, and Maintenance

Data storage has two status dimensions: state and status. State reflects administrative operations such as enabling, disabling, deletion, and maintenance. Status reflects the connection between data storage and hosts.

State

Data storage includes the following states:

- Enabled: Virtual machines can be created on the data storage.
- Disabled: Virtual machines and snapshots cannot be created on the data storage.
- Deleting: Data storage enters this state during deletion. No other operations are supported in this state.
- Maintenance Mode: Administrators can perform maintenance on data storage in this mode. After data storage enters maintenance mode, disks on the data storage become unavailable, and VM-related resources on the data storage are stopped.

Status

Status reflects the network connection between data storage and hosts. It includes Connecting, Connected, and Disconnected. When the network is unstable, the status of data storage might change. In the platform, reconnecting data storage can change its status.

6.6 Configuration

6.6.1 Modify Local Storage Configuration

After local storage is added, you can edit its name and description, adjust the data storage overcommit ratio, or modify the disk preallocation policy as needed.

Common Entries

You can modify local storage configuration from the following entries:

- Go to the target local storage page and click **Actions** > **Edit Name and Description**.
- Go to the target local storage page, click **Actions** > **Modify Advanced Settings**.

Common Operations

Operation	Description
Edit Name and Description	Modify the name and description of local storage.
Modify Data Storage Overcommit Ratio	Controls the allocatable data storage space for virtual machines. Calculation formula: `Data storage allocatable capacity = [(actual capacity - reserved capacity) ×`

Operation	Description
	`overcommit ratio] - (threshold capacity + sum of all volumes allocated to virtual machines + snapshots + image cache + migration cache)`
Modify Disk Preallocation Policy	Sets the disk preallocation policy for data storage. Default: **none**. • **none**: Does not use a preallocation policy. With this policy, files are dynamically allocated the required space as needed. • **metadata**: Preallocates only the space required for file metadata and does not allocate space for data. With this policy, storage space required for the volume is dynamically allocated as data is written. • **falloc**: Preallocates the required space for the file but does not immediately erase data on the corresponding physical device. Erasure is performed when the virtual machine writes data to the volume for the first time. • **full**: Preallocates the required space for the file.

6.6.2 Modify NFS Storage Configuration

After NFS storage is added, you can edit its name and description, adjust the data storage overcommit ratio, or modify the NFS storage disk preallocation policy as needed.

Common Entries

You can modify NFS storage configuration from the following entries:

- Go to the target NFS storage page and click **Actions** > **Edit Name and Description**.
- Go to the target NFS storage page, click **Actions** > **Modify Advanced Settings**.

Common Operations

Operation	Description
Edit Name and Description	Modify the name and description of NFS storage.
Modify Data Storage Overcommit Ratio	Controls the allocatable data storage space for virtual machines.

Operation	Description
	Calculation formula: `Data storage allocatable capacity = [(actual capacity - reserved capacity) × overcommit ratio] - (threshold capacity + sum of all volumes allocated to virtual machines + snapshots + image cache + migration cache)`
Modify NFS Storage Disk Preallocation Policy	Sets the disk preallocation policy for NFS storage. Default: **none**. • **none**: Does not use a preallocation policy. With this policy, files are dynamically allocated the required space as needed. • **metadata**: Preallocates only the space required for file metadata and does not allocate space for data. With this policy, storage space required for the volume is dynamically allocated as data is written. • **falloc**: Preallocates the required space for the file but does not immediately erase data on the corresponding physical device. Erasure is performed when the virtual machine writes data to the volume for the first time. • **full**: Preallocates the required space for the file.

6.6.3 Modify SAN Storage Configuration

After SAN storage is added, you can edit its name and description, modify the disk preallocation policy, adjust the SAN storage allocation strategy, or modify the data storage overcommit ratio as needed. For SAN storage expansion, see *Expand SAN Storage*.

Common Entries

You can modify SAN storage configuration from the following entries:

- Go to the target SAN storage page and click **Actions** > **Edit Name and Description**.
- Go to the target SAN storage page, click **Actions** > **Modify Advanced Settings**.

Common Operations

Operation	Description
Edit Name and Description	Modify the name and description of SAN storage.

Operation	Description
Modify Disk Preallocation Policy	Sets the disk preallocation policy for SAN storage. Default: **metadata**. • **metadata**: Preallocates only the space required for volume metadata and does not allocate space for data. With this policy, storage space required for the volume is dynamically allocated as data is written. • **none**: Does not use a preallocation policy.
Modify SAN Storage Allocation Strategy	Sets the placement strategy for volumes and snapshots on SAN storage LUNs. Default: Allocate by device name. Available policies include allocating by device name, creating volumes on the LUN with the most remaining capacity, and creating volumes on the LUN with the fewest volumes.
Modify Data Storage Overcommit Ratio	Controls the allocatable data storage space for virtual machines. Calculation formula: `Data storage allocatable capacity = [(actual capacity - reserved capacity) × overcommit ratio] - (threshold capacity + sum of all volumes allocated to virtual machines + snapshots + image cache + migration cache)`

6.6.4 Distributed Storage Configuration Management

After distributed storage is added, you can edit its name and description, adjust the data storage overcommit ratio, update the password for ZHPS distributed storage, manage monitor nodes, or manage storage pools as needed.

Common Entries

You can manage distributed storage configuration from the following entries:

- Go to the target distributed storage page and click **Actions** > **Edit Name and Description**.
- Go to the target distributed storage page, click **Actions** > **Modify Advanced Settings**.
- Go to the target ZHPS distributed storage page and click **Actions** > **Update Password**.
- Go to the target distributed storage page and open the **Monitoring Node** tab.
- Go to the target distributed storage page and open the **Storage Pool** tab.

Common Operations

Operation	Description
Edit Name and Description	Modify the name and description of distributed storage.
Modify Data Storage Overcommit Ratio	Controls the allocatable data storage space for virtual machines. Calculation formula: `Data storage allocatable capacity = [(actual capacity - reserved capacity) × overcommit ratio] - (threshold capacity + sum of all volumes allocated to virtual machines + snapshots + image cache + migration cache)`
Update ZHPS Distributed Storage Password	Enter the **New Password** used by the storage platform, enter it again in **Confirm Password**, and then click **OK**.  **Note:** After the password is updated, the platform automatically reconnects the data storage. If the password does not match, the data storage cannot be reconnected. This operation updates only the connection password stored on the platform and does not change the password on the storage platform.
Add Monitor Node	On the **Monitoring Node** tab, click **Add Monitor Node**, specify monitoring node management IP, SSH port, username, and password.
Modify Monitor Node Configuration	On the **Monitoring Node** tab, click **Actions** for the target monitoring node and modify the SSH username, SSH password, SSH port, and Mon port as needed.
Delete Monitor Node	On the **Monitoring Node** tab, click **Actions** > **Delete Monitor Node** for the target monitor node.  **Note:** Deleting monitoring nodes might cause the distributed storage cluster to lose connectivity. Proceed with caution.

Operation	Description
Add Storage Pool	On the **Storage Pool** tab, click **Add Storage Pool**, specify the pool name and display name.
Modify Storage Pool Display Name	On the **Storage Pool** tab, click **Actions** > **Set Display Name**.
Delete Storage Pool	On the **Storage Pool** tab, click **Actions** > **Delete**. **Note:** • The image cache pool cannot be deleted. • A storage pool cannot be deleted if it contains data. • At least one storage pool must be retained.

6.7 Capacity and Space Management

6.7.1 Expand SAN Storage

When SAN storage capacity can no longer meet business requirements, you can add new LUNs or expand existing LUNs.

Prerequisites

Before expanding SAN storage, expand the LUNs of iSCSI, FC, or NVMe storage in the physical environment.

Procedure

1. Expand the LUNs of iSCSI storage, FC storage, or NVMe storage in the physical environment.

   After LUN expansion succeeds, return to the platform. On the **Data Center** > **Storage Target** tab, you can view block device usage or manually click the refresh button to obtain the latest storage information.
2. Navigate to the target SAN storage page.
3. Go to the **Shared Block** tab.
4. Click **Add** and select the newly expanded LUN.

   Note:
   If data exists on the LUN, adding the LUN or attaching SAN storage might fail. You can select **Cleanup Block Device** on the **Add Shared Block** page.

What's next

To obtain the latest capacity information of a LUN, select the target LUN on the **Shared Block** tab and click **Refresh Capacity**.

6.7.2 Clean Up Data Storage

You can clean up original data retained in data storage after cross-storage migration as needed.

Procedure

On the **Cleanup Data** tab of the target data storage details page, click **Actions** > **Cleanup**.

Note:
Make sure the migrated data is intact. Original data cannot be recovered after cleanup. Proceed with caution.

6.8 Delete Data Storage

Prerequisites

Before deleting data storage, detach it from the cluster. Otherwise, it cannot be deleted.

- Delete a single data storage resource: Navigate to the target data storage page and click **Actions** > **Delete**.
- Delete data storage in batch: Navigate to **Data Center** > **Data Storage**, select the data storage resources to delete, and click **Actions** > **Delete**.

What's next

Note:
During deletion, all resources on the selected data storage, including virtual machines, volumes, and snapshots, are deleted. Proceed with caution.

6.9 Distributed Storage Service Deployment and Takeover

The distributed storage service is used to deploy or take over distributed storage systems in NexaSphere. You can initialize a distributed storage cluster through the quick deployment workflow, or take over an external distributed storage system that has already been deployed or is already running.

After deployment or takeover, the platform can centrally manage and use distributed storage resources, and provide information such as storage overview, node details, performance metrics,

and disk health status. You can add distributed storage as data storage for virtual machines based on business requirements.

6.9.1 Deploy Distributed Storage in 3 Steps

6.9.1.1 Step 1: Upload the Installation Package

Prerequisites

- You have prepared the installation package of the distributed storage software.
- The platform management node has been added to the platform as a host.
- At least two distributed storage nodes have been added to the platform as hosts, and the storage network has been configured for the nodes by using Kernel adapters from hosts or distributed port groups. For more information about configuring the storage network, see *Create a Kernel Adapter*.
- To ensure that distributed storage can enable management node HA together with the platform, configure management node HA in MN Ops before installing distributed storage. For more information about configuring management node HA, see *Set Up with the UI*.

Procedure

1. In the navigation pane, choose **Storage Service** > **Storage Overview**.
2. On the **Storage Overview** page, click **Upload**.
3. In the **Upload Installation Package** dialog, configure the following options:

Option	Description
Server IP	Displays the IP address of the current management node. In a dual management node environment, this field displays the IP address of the node that hosts the VIP.
Storage Path	Specify a directory path on the server to store uploaded installation packages and their extracted files. **Note:** • Make sure this directory has sufficient available space. Otherwise, installation might fail.

Option	Description
	• Do not use system directories such as `/`, `/dev`, `/proc`, `/sys`, `/usr/bin`, `/bin`, or `/opt`. Using system directories might cause server exceptions.
Upload By	Select how to upload the installation package. Options include URL and Local File. Interrupted local uploads can be resumed.  **Note:** • The system automatically retries recoverable upload errors. If an upload is paused due to a network interruption and the page remains open, the upload resumes automatically after the network recovers. • If the browser retains permission to access the original local file after the page is refreshed or you sign in again, the system automatically resumes an upload that was paused due to an error. If the browser prompts you to allow access to previously selected files or folders, grant access as needed. • If the browser does not have permission to access the file or the permission has expired, the upload cannot resume automatically. Click **Pending Uploads** at the top of the page. In the **Pending Upload Tasks** panel, locate the task, click **Select to Continue**, and reselect the same local file. • A manually paused task does not resume automatically. In the task details, click **Continue**. • If the local file has been modified, moved, or deleted, the original upload task cannot resume. Start a new upload. • If an upload error occurs, the platform automatically retries the upload up to five times. If all five retry attempts fail, the upload task enters the Failed state and cannot be resumed. Start a new upload. • Click **View Upload History** to view local file upload tasks.

4. Click **OK**.

6.9.1.2 Step 2: Deploy the Management Service

Procedure

1. On the **Storage Overview** page, click **Install**.
2. In the **Deploy Management Service** dialog, configure the following options:

Option	Description
VIP	Displays the management node IP address in a single management node environment, or the VIP address in a dual management node environment.
MN IP	Displays the IP address of the single management node or dual management nodes. Enter the SSH port and password of the management node.
Database Password	If left empty, the initial database password is used by default. If the database password has been changed, enter the new password.

3. Review the configuration and click **OK**.

6.9.1.3 Step 3: Initialize Distributed Storage

Prerequisites

Before initialization, configure the storage network for distributed storage nodes by using Kernel adapters from hosts or distributed port groups. For more information about configuring the storage network, see *Create a Kernel Adapter*.

Procedure

1. On the **Storage Overview** page, click **Initialize**.
2. In the **Initialize Distributed Storage** dialog, complete initialization and storage configuration.
3. For **Initialization**, configure the following options:

Option	Description
Storage MN IP	Select the management node IP address of the distributed storage.
Cluster	Displays the cluster location.
Admin Network	Enter the admin network CIDR to manage and configure the storage cluster.

Option	Description
Cluster Network	Enter the cluster network CIDR to monitor data disks and synchronize replicas among nodes in the storage cluster.
Public Network	Enter the public network CIDR for the storage cluster to provide external services.
Time Sync Server IP	This parameter is displayed based on the storage service type detected by the system. Enter the time synchronization server IP address to ensure time synchronization between nodes in the storage cluster.
Gateway Network	This parameter is displayed based on the storage service type detected by the system. Enter the gateway network so that user services can access storage resources through this network.

4. Review the configuration and click **Next**.
5. For **Storage Configuration**, click **Add** to add distributed storage nodes.
6. Configure manager and monitor roles.

   For data security, we recommend assigning manager (Mgr) and monitor roles (Mon) to three nodes.
7. Specify hostnames for the nodes.

   Note:

   - Length: 2 to 60 characters.
   - Allowed characters: uppercase and lowercase letters, numbers, and hyphens (-).
   - The hostname cannot contain consecutive hyphens or start or end with a hyphen.
   - When nodes are added, suffixes such as `-1/-2/-3` are automatically appended to node hostnames to ensure uniqueness.
8. Review the configuration and click **OK**.

What's next

Starting from NexaSphere 1.10.20, if you need to access the distributed storage management interface separately:

- For nSDS-X v6.4.200.1, the default login credentials are `admin/Admin@123`.

- For nSDS-X v5, the default login credentials remain `admin/password`.

6.9.2 Take Over Existing Distributed Storage

6.9.2.1 Take Over Distributed Storage

After storage package installation, network configuration, and system initialization are complete on the server, or when distributed storage is already running and in use, you can use this method to take over distributed storage.

Procedure

1. In the navigation pane, choose **Storage Service** > **Storage Overview**.
2. On the **Storage Overview** page, click **Take Over Existing Storage**.
3. In the **Take Over Distributed Storage** dialog, configure the following options:

Option	Description
Storage MN IP	Enter the management node IP address of the distributed storage to be taken over.
Storage Service Check	Click **Check** to check whether storage services exist and whether the connection can be established successfully.
Username	Enter the username for logging in to the distributed storage platform.
Password	Enter the login password.
Access Token	Enter this parameter as needed based on the detected storage service type.

4. Review the configuration and click **OK**.

What's next

Starting from NexaSphere 1.10.20, if you need to access the distributed storage management interface separately:

- For nSDS-X v6.4.200.1, the default login credentials are `admin/Admin@123`.
- For nSDS-X v5, the default login credentials remain `admin/password`.

6.9.2.2 Cancel Distributed Storage Takeover

Canceling distributed storage takeover does not affect stored data, but the platform can no longer manage the distributed storage system.

Procedure

1. In the navigation pane, choose **Storage Service** > **Storage Overview**.
2. On the **Storage Overview** page, click **Cancel Takeover**.
3. In the confirmation dialog, read the risk warning and click **OK** after you acknowledge the risk.

7 Networking

7.1 Overview

NexaSphere organizes virtual networking through distributed switches, uplinks, distributed port groups, and Kernel adapters to provide consistent network connectivity for hosts, virtual machines, and internal platform traffic. After initializing basic resources, you can configure network resources based on service network, storage network, and access control requirements.

Distributed switches provide consistent virtual switching across hosts. Uplinks connect host physical NICs to distributed switches. Distributed port groups carry network configuration for VM NICs and Kernel adapters. Kernel adapters identify host-side management network or storage network traffic.

In addition to basic network resources, NexaSphere provides network topology and security groups. Network topology shows the connection relationships among distributed switches, clusters, hosts, distributed port groups, and virtual machines. Security groups control ingress and egress traffic for VM NICs.

7.2 Network Model and Resources

NexaSphere networking organizes virtual network resources around distributed switches. Hosts connect to distributed switches through uplinks. VM NICs and Kernel adapters use the corresponding networks through distributed port groups. Security groups provide traffic access control on VM NICs.

Resource	Description
Distributed switch	Provides unified virtual switching for clusters and hosts in a data center. It is the central resource for distributed port groups, uplinks, and network topology.
Uplink	Connects host physical NICs to a distributed switch. Uplink configuration includes aggregation mode, hash policy, and host NIC selection, and determines how hosts connect to the corresponding virtual network.

Resource	Description
Distributed port group	Provides network configuration for VM NICs and Kernel adapters. A distributed port group can include VLAN, DHCP, IP address management, network segment, and DNS configuration.
Kernel adapter	Uses network labels to identify host-side traffic. After the first host is added to a cluster, NexaSphere creates a default Kernel adapter for each host to obtain and display the host management network. You can also create Kernel adapters for storage network services.
Security group	Controls VM NIC traffic through ingress and egress security group rules. A VM NIC can join multiple security groups, and traffic is matched against rules based on security group priority.

Default Network Resources

After the first host is added to a cluster, NexaSphere automatically creates a default distributed switch, a default distributed port group, and a default Kernel adapter based on the host configuration. These resources are used for centralized management of the host management network. You can later reuse the default distributed switch to create distributed port groups or create custom distributed switches based on your network plan.

7.3 Configure Distributed Switches

7.3.1 Create a Distributed Switch

Procedure

1. In the navigation pane, choose **Inventory** > **Network Resource**.
2. In the resource tree, right-click the target data center and select **New Distributed Switch**.
3. In the **New Distributed Switch** dialog, complete basic information configuration.

Option	Description
Name	Enter a name for the distributed switch.
Description	Enter a brief description for the distributed switch.
Data Center	Displays the data center where the distributed switch resides.

Option	Description
Cluster	Select the cluster to attach to the distributed switch. Virtualization clusters and bare metal clusters are supported.

4. Complete network configuration as needed.

   Network configuration is required only when you select a virtualization cluster. Select an addition method based on your host NIC plan, and then configure the corresponding options.

   Add Individually or Batch Bonding

Option	Description
Addition Method	Select **Add Individually** or **Batch Bonding**.
Uplink Name	Set a name for the bonded host physical NICs that connect to physical switches. • By default, the name uses the `Uplink+suffix` format, and the suffix increments as `1/2/3`. When the number of uplinks is 10 or more, the default name uses the `Up+suffix` format. • A custom name must be 1 to 10 characters long, can contain only letters, numbers, hyphens, and underscores, and cannot start with a number.
Bond Mode	Select a bond mode for physical NICs. • In LACP mode (mode 4), bonded NICs share the same speed and duplex settings, and network traffic is distributed across the NICs. • In active-backup mode (mode 1), the active NIC handles network traffic during normal operation and the backup NIC takes over automatically if the active NIC fails. • Both modes support 1 to 8 physical NICs, and 2 or more NICs are recommended.
Hash Policy	When LACP mode is selected, set a hash policy to determine the network traffic egress.

Option	Description
	• **layer 2+3**: Picks out a NIC port to send data packets based on the hash computation on the source MAC address, destination MAC address, and IP address. • **layer 3+4**: Picks out a NIC port to send data packets based on the hash computation on the IP address and port. TCP/IP stacks are supported. • **layer 2**: Picks out a NIC port to send data packets based on the hash computation on the source MAC address and destination MAC address.
Host NIC Ports	Select the physical ports on the host to bond. • When creating bonds for each individual host, the selected ports on that host must operate at the same speed. • When creating bonds for multiple hosts in bulk, you can only select ports that have the same speed.

Specify the Same Port

Option	Description
Addition Method	Select **Specify the Same Port**.
Port Type	Select the port type, including bond interfaces and non-bonded port.
Bond/Host NIC Ports	Select a NIC to add all hosts in the cluster that have the same NIC.

5. Complete distributed port group configuration as needed.

   By default, a distributed port group is created. You can choose whether to create a distributed port group on the distributed switch.

Option	Description
Name	Enter a name for the distributed port group.

Option	Description
VLAN Type	Select a VLAN type. If you select standard VLAN, enter the VLAN ID.
DHCP Service	Select whether to enable the service that automatically assigns IP addresses to internal platform resources.
DHCP Service IP	Set the IP address used by the DHCP service.
DNS	Set the DNS resolution service for the distributed port group.

IP Address Management

After IP address management is enabled, you can add network ranges to the distributed port group. IP addresses in the network ranges can be assigned through DHCP to resources that use this network.

IPv4 Address Type

Option	Description
IP Allocation Policy	After the DHCP service is enabled, the following allocation policies are supported: • Random allocation: The system randomly assigns IP addresses in the added network range. • Allocate in Order: The system assigns all available IP addresses in the added network range in ascending order. Released IP addresses are assigned in the next allocation. Example: Assume that the network range is *192.168.0.101 ~ 192.168.0.120*, within which *192.168.0.101 ~ 192.168.0.108* are allocated. If *192.168.0.106* is released, it will be assigned first in the next allocation. • Allocate in Cycle: The system assigns IP addresses in the added network range in ascending order. Addresses released during allocation are assigned after all currently available IP addresses have been assigned once.

Option	Description
	Example: Assume that the network range is *192.168.0.101 ~ 192.168.0.120*, within which *192.168.0.101 ~ 192.168.0.108* are allocated. If *192.168.0.106* is released, it will be assigned after *192.168.0.120* is used.
Network Range Method	The following network range methods are supported: • IP Range: Enter the start IP address, end IP address, subnet mask, and gateway of the network range. **Note:** Do not include the gateway, broadcast address, or network address in the added IP range. • CIDR: Enter the CIDR block and gateway. The gateway can be the first or last address in the CIDR block. If the gateway is left empty, the first address in the CIDR block is used by default.

IPv6 Address Type

Option	Description
Network Range Method	The following network range methods are supported: • IP Range: Enter the start IP address, end IP address, prefix length, and gateway of the network range. • CIDR: Enter the CIDR block.
IP Allocation Mode	Select the IPv6 address allocation mode: • **Stateful-DHCP**: Stateful DHCP configuration. The interface address and other parameters are all configured through DHCP. The IP range method supports only this allocation mode. • **Stateless-DHCP**: Stateless DHCP configuration. The interface address is derived from the prefix in router advertisements

Option	Description
	and the interface MAC address, and other parameters are configured through DHCP. • **SLAAC**: Stateless address autoconfiguration. The interface address is derived from the prefix in router advertisements, and other parameters are also included in router advertisements.

6. After confirming that the configuration is correct, click **OK**.

Note:

- After a single NIC or bonded NIC is added to the distributed switch, bridge names created by distributed port groups use the `br_dvs` prefix.
- If the management network address is on a NIC or sub-interface before the host is added, the management network address is moved to the `br_dvs{ID}_{VLAN ID}` bridge after the host is added. After the host is deleted, the management network address moves back to the NIC or sub-interface.
- If the management network address is on a `br` bridge before the host is added, for example a user-defined `br_{name}_{VLAN ID}` bridge, the bridge name remains unchanged after the host is added. When the management network and service network share NICs, if the VLAN of the new distributed port group is different from the management network VLAN, the bridge created by the distributed port group still uses the `br_dvs{ID}_{VLAN ID}` format, but the NIC that carries the management network address is not affected.

7.3.2 Attach or Detach Clusters

Prerequisites

The default distributed switch supports attaching or detaching only bare metal clusters.

Procedure

1. In the navigation pane, choose **Inventory** > **Network Resource**.
2. In the resource tree, select the target distributed switch.
3. On the target distributed switch details page, click the **Cluster** tab.
4. Attach clusters as needed.
   a) Click **Attach Cluster**.

b) In the **Attach Cluster** dialog, select the target cluster, bond configuration method, and host NIC.

When attaching a cluster, you cannot specify the bond mode. The bond mode is inherited from the distributed switch.

c) Click **OK**.

5. Detach clusters as needed.

a) Click **Detach Cluster**.

b) In the **Detach Cluster** dialog, select the target cluster.

c) Click **OK**.

Note:

After a cluster is detached, the corresponding VM NICs are detached. Proceed with caution.

7.3.3 Delete a Distributed Switch

Prerequisites

- You cannot delete the default distributed switch while it still has joined hosts.
- You cannot delete a distributed switch if a distributed port group on the distributed switch is referenced by a VM memory snapshot.
- You cannot delete a distributed switch if a distributed port group on the distributed switch is associated with a Kernel adapter.

Procedure

1. In the navigation pane, choose **Inventory** > **Network Resource**.
2. In the resource tree, select the target distributed switch.
3. On the target distributed switch details page, click **Actions** > **Delete**.

Note:

Deleting a distributed switch deletes its distributed port groups and detaches the corresponding VM NICs. Proceed with caution.

7.4 Manage Uplinks

7.4.1 Modify Uplink Configuration

Procedure

1. In the navigation pane, choose **Inventory** > **Network Resource**.
2. In the resource tree, select the target distributed switch.
3. On the target distributed switch details page, click the **Uplink** tab.
4. Click **Modify Configuration**.
5. In the **Modify Uplink Mode** dialog, modify the bond mode and hash policy as needed.

   Note:

   - You can modify the bond mode when hosts have joined the default distributed switch only if the default distributed switch uses a single NIC.
   - You can modify the bond mode of a default distributed switch when no hosts have joined it.
6. Click **OK**.

7.4.2 Manage Joined Hosts

Procedure

1. In the navigation pane, choose **Inventory** > **Network Resource**.
2. In the resource tree, select the target distributed switch.
3. On the target distributed switch details page, click the **Uplink** tab.
4. In the **Joined Host** list, select the target host.

   a) To add physical NICs to the bond, click **Actions** > **Add Physical Port**.

   Note:

   1. You can add only NICs that are not bonded.
   2. New NICs must have the same speed as existing NICs in the bond.
   3. A bond supports up to 8 physical NICs.

   b) To remove physical NICs from the bond, click **Actions** > **Remove Physical Port**.

   Note:

   You cannot remove all physical NICs. At least one physical NIC must remain.

   c) To remove the host's entire uplink, click **Actions** > **Disconnect Host Uplink**.

   Note:

   1. Disconnecting a host uplink also deletes the bond on the host. Proceed with caution.

2. Disconnecting a host uplink cascades to detach VM NICs on the host. Proceed with caution.
3. The default distributed switch does not support disconnecting host uplinks.

7.4.3 Configure Uplinks for Hosts

Procedure

1. In the navigation pane, choose **Inventory** > **Network Resource**.
2. In the resource tree, select the target distributed switch.
3. On the target distributed switch details page, click the **Uplink** tab.
4. In the **Unjoined Host** list, select the target host.
   a) If the new host bond configuration matches the uplink configuration of the distributed switch associated with the host's cluster, the host automatically joins the uplink.
   b) If the new host bond configuration does not match the uplink configuration of the distributed switch associated with the host's cluster, click **Actions** > **Join Uplink**.
   After the host joins the uplink, the host bond configuration is adjusted to match the switch uplink configuration.

7.5 Configure Distributed Port Groups

7.5.1 Create a Distributed Port Group

Procedure

1. In the navigation pane, choose **Inventory** > **Network Resource**.
2. In the resource tree, right-click the target distributed switch and select **New Distributed Port Group**.
3. In the **New Distributed Port Group** dialog, complete basic configuration.

Option	Description
Name	Enter a name for the distributed port group.
Description	Enter a brief description for the distributed port group.
Distributed Switch	Select the distributed switch for the distributed port group.
VLAN Type	Select a VLAN type. If you select standard VLAN, enter the VLAN ID.

Option	Description
DHCP Service	Select whether to enable the service that automatically assigns IP addresses to internal platform resources.
DHCP Service IP	Set the IP address used by the DHCP service.
DNS	Set the DNS resolution service for the distributed port group.

4. Enable IP address management and configure network ranges as needed.

   After IP address management is enabled, you can add network ranges to the distributed port group. IP addresses in the network ranges can be assigned through DHCP to resources that use this network.

   IPv4 Address Type

Option	Description
IP Allocation Policy	After the DHCP service is enabled, the following allocation policies are supported: • Random allocation: The system randomly assigns IP addresses in the added network range. • Allocate in Order: The system assigns all available IP addresses in the added network range in ascending order. Released IP addresses are assigned in the next allocation. Example: Assume that the network range is *192.168.0.101 ~ 192.168.0.120*, within which *192.168.0.101 ~ 192.168.0.108* are allocated. If *192.168.0.106* is released, it will be assigned first in the next allocation. • Allocate in Cycle: The system assigns IP addresses in the added network range in ascending order. Addresses released during allocation are assigned after all currently available IP addresses have been assigned once. Example: Assume that the network range is *192.168.0.101 ~ 192.168.0.120*, within which *192.168.0.101 ~ 192.168.0.108* are allocated. If *192.168.0.106* is released, it will be assigned after *192.168.0.120* is used.

Option	Description
Network Range Method	The following network range methods are supported: • IP Range: Enter the start IP address, end IP address, subnet mask, and gateway of the network range. **Note:** Do not include the gateway, broadcast address, or network address in the added IP range. • CIDR: Enter the CIDR block and gateway. The gateway can be the first or last address in the CIDR block. If the gateway is left empty, the first address in the CIDR block is used by default.

IPv6 Address Type

Option	Description
Network Range Method	The following network range methods are supported: • IP Range: Enter the start IP address, end IP address, prefix length, and gateway of the network range. • CIDR: Enter the CIDR block.
IP Allocation Mode	Select the IPv6 address allocation mode: • **Stateful-DHCP**: Stateful DHCP configuration. The interface address and other parameters are all configured through DHCP. The IP range method supports only this allocation mode. • **Stateless-DHCP**: Stateless DHCP configuration. The interface address is derived from the prefix in router advertisements and the interface MAC address, and other parameters are configured through DHCP. • **SLAAC**: Stateless address autoconfiguration. The interface address is derived from the prefix in router advertisements, and other parameters are also included in router advertisements.

5. After confirming that the configuration is correct, click **OK**.

7.5.2 Modify Distributed Port Group Configuration

Procedure

1. In the navigation pane, choose **Inventory** > **Network Resource**.
2. In the resource tree, select the target distributed port group.
3. Modify basic configuration as needed.
   a) Click **Actions** > **Modify Configuration**.
   b) In the **Modify Configuration** dialog, modify the name, description, VLAN ID, MTU, and DHCP service as needed.
   c) Click **OK**.
4. Add or delete network ranges as needed.

   Click the **Network Range** tab, and then perform the corresponding operation as needed.

   Note:

   - Deleting a network range detaches the VM NICs that use the network range. Proceed with caution.
   - If the distributed port group already has a DHCP service IP and the DHCP service IP is in the selected network range, deleting all network ranges from the distributed port group releases the DHCP service IP. If only some network ranges are deleted, the DHCP service IP is retained.
5. Modify DNS configuration as needed.

   Click the **DNS** tab, and then perform the corresponding operation as needed.

7.5.3 Delete a Distributed Port Group

Prerequisites

- You cannot delete the default distributed port group.
- You cannot delete a distributed port group if it is referenced by a VM memory snapshot.
- You cannot delete a distributed port group if it is associated with a Kernel adapter.

Procedure

1. In the navigation pane, choose **Inventory** > **Network Resource**.
2. In the resource tree, select the target distributed port group.
3. On the target distributed port group details page, click **Actions** > **Delete**.

Note:

Deleting a distributed port group detaches the VM NICs that are using this port group. Proceed with caution.

4. After acknowledging the risk, click **OK**.

7.6 Configure Kernel Adapters

7.6.1 Create a Kernel Adapter

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. In the resource tree, select the target host.
3. On the target host details page, click the **Kernel Adapter** tab.
4. On the **Kernel Adapter** tab, click **New Kernel Adapter**.
5. In the **New Kernel Adapter** dialog, complete configuration.

Option	Description
Name	Enter a name for the Kernel adapter.
Description	Enter a brief description for the Kernel adapter.
Network Service	Displays **Storage** by default, indicating that the Kernel adapter handles storage network traffic.
Distributed Port Group	Select a distributed port group.
IPv4 Address	Assign an IPv4 address to the Kernel adapter.
Netmask	Specify a netmask.

6. After confirming that the configuration is correct, click **OK**.

7.6.2 Modify Kernel Adapter Configuration

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. In the resource tree, select the target host.
3. On the target host details page, click the **Kernel Adapter** tab.

4. On the **Kernel Adapter** tab, select a Kernel adapter from the list, and then click **Actions** > **Modify Configuration**.
5. In the **Modify Configuration** dialog, make changes as needed.
   - For the default Kernel adapter, you can modify only the name, description, and whether the **Storage** network service is selected. When **Storage** is selected, the storage network shares the management network.
   - For other Kernel adapters, you can modify the name, description, IPv4 address, and netmask.
6. Click **OK**.

7.6.3 Delete a Kernel Adapter

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. In the resource tree, select the target host.
3. On the target host details page, click the **Kernel Adapter** tab.
4. On the **Kernel Adapter** tab, select a Kernel adapter from the list, and then click **Actions** > **Delete**.

   Note:

   - You cannot delete the default Kernel adapter.
   - Deleting a Kernel adapter releases the associated IP address. Storage services that depend on this IP address are interrupted. Proceed with caution.
5. After entering the confirmation information, click **OK**.

7.7 View Network Topology

Context

Network topology is centered around a distributed switch and shows the network relationships among clusters, hosts, distributed port groups, and virtual machines associated with the distributed switch.

Procedure

1. In the navigation pane, choose **Inventory** > **Network Resource**.
2. In the resource tree, select the target distributed switch.

3. On the target distributed switch details page, click the **Network Topology** tab.
4. View the resource connection relationships displayed in the network topology.

7.7.1 Supported Network Topology Operations

When viewing network topology, you can use the following operations to adjust the topology view or view resource relationships.

Operation	Description
Refresh	Displays the latest network topology.
Zoom In / Zoom Out	Zooms in or out on the network topology.
Default Position	Returns to the origin of the topology canvas.
Export	Exports the network topology as a PNG image.
Hide / Show Virtual Machines	Hides or shows virtual machines in the network topology.
Full Screen	Displays the network topology in full screen.
Highlight	Selects a resource and highlights the resource and its associated resources.
Hover	Displays resource information when you hover over a resource, and supports opening the resource details.
Search	Searches topology resources by resource name or UUID.

7.8 Manage Security Groups

A security group is a collection of security rules that control ingress and egress traffic for VM NICs . A VM NIC can join multiple security groups, and the security group priority determines the rule matching order.

Rule Types

Security groups use security rules in the group to control traffic to and from NICs. A security group can contain system rules and custom rules.

- System rules: After a security group is created, the system provides the following rules by default:
  - Intra-group communication rule: Allows NICs in the same security group to communicate with each other by default. This rule has a higher priority than all custom rules. It cannot be modified or deleted and can only be disabled.

    - Intra-/inter-group communication rule: Allows NICs in the security group to access NICs outside the security group by default. NICs outside the security group are not allowed to access NICs inside the security group by default. This rule supports modifying the default intra-group and inter-group access behavior for individual VM NICs.
- Custom rules: Rules that users add to the security group.

Rule Elements

Element	Description
Direction	• Ingress rule: Controls traffic that enters the NIC from outside and mainly controls the traffic source. • Egress rule: Controls traffic sent out from the NIC and mainly controls the traffic destination.
Target	• Source: Corresponds to ingress rules and supports IP addresses /ranges or security groups as sources. An ingress rule allows or denies traffic from the specified IP addresses/ranges or security groups. • Destination: Corresponds to egress rules and supports IP addresses/ranges or security groups as destinations. An egress rule allows or denies NICs in the current group from accessing the target IP addresses/ranges or NICs in the target security groups.
Action	The following actions can be applied to traffic that matches rule conditions: • Allow: Allows network request traffic to flow into or out of the NIC. • Deny: Does not allow network request traffic to flow into or out of the NIC. By default, if traffic to or from the NIC does not match any custom rule, ingress traffic is denied and egress traffic is allowed.
Protocol and Port	The packet protocol and corresponding port targeted by a rule include the following types: • ALL: Covers all protocol types. Ports cannot be specified.

Element	Description
	• TCP: Supports ports 1-65535. • UDP: Supports ports 1-65535. • ICMP: Does not support specifying a port.
Priority	The matching order of a security group rule relative to other rules. The supported priority range is 1-100. A larger value indicates a lower priority.

7.8.1 Create Security Groups and Rules

Context

Before a virtual machine can use a security group, you need to create the security group, add rules to the security group, and associate the security group with VM NICs.

Procedure

1. Create a security group.
   a) In the target data center, choose **Network** > **Security Group** > **Create Security Group**, and configure the following options.

Option	Description
Name	Enter a name for the security group.
Description	Enter a description for the security group.

   b) Click **OK**.
2. Add rules to the security group.
   a) Click the security group name to enter its details page.
   b) On the security group **Overview** tab, select the **Ingress Rule** or **Egress Rule** tab based on the rule direction.
   c) Click **Add Rule**.
   d) In the **Add Ingress Rule/Add Egress Rule** dialog, configure the following options:

Option	Description
Type	The traffic direction controlled by the rule, displayed as ingress or egress.

Option	Description
Priority	The effective priority of the rule. The value automatically increases by 1 for each new rule. A larger value indicates a lower priority.
Policy	Specify a policy for the rule. Options include Allow and Reject.
IP Address Type	Specify IPv4 or IPv6.
Protocol	The communication protocol targeted by the rule. ALL, TCP, UDP, and ICMP are supported.
Port	When TCP or UDP is selected, specify the ports targeted by the rule. • Use the `Start Port-End Port` format for a port range. • Separate multiple ports or port ranges with commas. You can specify up to 10 ports or port ranges.
Source	Set this option when adding an ingress rule. It indicates whether to allow or deny access from the specified IP addresses/ranges or security groups. • When specifying IP addresses/ranges, you can enter an IP range in the `Start IP-End IP` format. • When specifying IP addresses/ranges, you can enter CIDR notation. If CIDR and other types of IP addresses are specified together, the CIDR mask must be 24 bits. If only CIDR is specified, there is no mask restriction. • Separate multiple IP addresses/ranges with commas.
Destination	Set this option when adding an egress rule. It indicates whether to allow or deny NICs in the group from accessing the specified IP addresses or security groups. • When specifying IP addresses/ranges, you can enter an IP range in the `Start IP-End IP` format. • When specifying IP addresses/ranges, you can enter CIDR notation. If CIDR and other types of IP addresses are

Option	Description
	specified together, the CIDR mask must be 24 bits. If only CIDR is specified, there is no mask restriction. • Separate multiple IP addresses/ranges with commas.
State	Specifies whether to enable the rule immediately after the security group is created. The rule is enabled by default. If it is disabled, NICs in the group do not match this rule until you manually enable it.
Description	Enter a description for the security group rule.

e) Review the configuration and click **OK**.

3. Import or export rules in batch as needed.
   - To add rules in batch, on the security group details page, click **Actions** > **Import Rule** and upload a CSV file. Imported rules do not affect existing rules. Their priorities are placed after existing rules by default, and the rules are disabled by default. After importing the rules, you can manually adjust their priorities and enable them. To ensure system compatibility, import files must be edited by using Microsoft Excel.
   - To reuse rules from one security group in another security group, export ingress rules, egress rules, or all rules, and then import them into the target security group.
4. Associate the security group with VM NICs.
   - To associate the security group with VM NICs in batch, go to the security group **VM NIC** tab, select the target NICs in batch, and select the network scope to which the security group applies and the target NICs.
   - To associate VM NICs with security groups in batch, go to the target virtual machine **Overview** tab and use **Modify Configuration** to associate security groups with VM NICs in batch. The smaller the value corresponding to a security group bound to a virtual machine, the higher the effective priority of the security group.

Note:
If the VM NIC IP address is empty on the platform, the default security group rule, namely the intra-group communication rule, does not take effect for the NIC. Configure security groups properly to avoid rule conflicts between security groups.

7.8.2 Operate Security Groups and Rules

Context

After creating security groups and rules and binding security groups to VM NICs, you can modify security group state, configuration, bindings, and rule priorities as needed.

- Enable or disable a security group.

  To modify the enabled or disabled state of a security group, use the **Enable** or **Disable** operation.

- Modify the name and description of a security group.

  In the operation column of the target security group, click **Edit Name and Description**, and then make changes as needed.

- Disassociate VM NICs.

  - To disassociate VM NICs from a security group in batch, go to the security group **VM NIC** page, select the target NICs, and click **Disassociate VM NIC**.
  - To disassociate security groups from VM NICs in batch, go to the target virtual machine **Overview** page and use **Modify Configuration** to unbind security groups from VM NICs in batch.

- Delete a security group.

  In the operation column of the target security group, click **Delete**.

   **Note:**
  Deleting a security group also deletes custom rules created for the security group. Proceed with caution.

- Enable or disable a security group rule.

  Go to the target security group **Overview** page, select the **Ingress Rule** or **Egress Rule** tab, and then use the **Enable** or **Disable** operation to modify the rule state. The system-default intra-group communication rule can only be disabled.

- Modify a security group rule.

  In the operation column of the target security group rule, click **Edit Rule**, and modify the priority, allow or deny policy, protocol and port, source or destination, state, and description as needed.

- Adjust security group rule priorities.

Click **Adjust Priority**, and drag rules to adjust their relative priority order. The priority range is 1-100. A larger value indicates a lower priority.

- Delete a security group rule.

  Select the target rule and click **Delete**.

8 Resource Management and Availability

8.1 Overview

NexaSphere provides resource allocation, dynamic resource scheduling, VM scheduling policies, and VM HA policies to help administrators control how resources are used in a virtualized environment and maintain service continuity when host loads change or VM-related resources fail.

Resource allocation and overcommit define the CPU, memory, and storage capacity that can be allocated to VMs. Administrators can plan VM resource usage based on CPU overcommit ratio, memory overcommit ratio, data storage overcommit ratio, and reserved capacity, and use the capacity calculation rules to understand the remaining allocatable capacity.

Cluster DRS monitors host loads in a cluster and provides migration suggestions or migrates VMs based on manual or automatic scheduling policies. VM scheduling policies define affinity or anti-affinity relationships between VMs, or between VMs and hosts. VM HA policies trigger automatic restart or migration when a VM stops unexpectedly or related resources fail.

8.2 Resource Allocation and Overcommit

Resource allocation describes the CPU, memory, and storage resources that are occupied or reserved by VMs. Resource overcommit allows physical resources to be virtualized at a specified ratio and then allocated to VMs.

CPU overcommit ratio, memory overcommit ratio, and data storage overcommit ratio together affect the upper limit of resources available to VMs. Host reserved memory reserves the memory required by the KVM host system. Data storage reserved capacity and threshold capacity help protect storage space. When planning resources, consider the physical total, reserved capacity, overcommit ratio, allocated capacity, and remaining allocatable capacity together.

Resource Allocation Terms

Term	Description
Physical total	Total physical CPU, physical memory, or data storage capacity that can participate in allocation.
Reserved capacity	Resource capacity reserved for host system operation or storage safety. Memory calculation first deducts host reserved memory, and

Term	Description
	storage calculation first deducts reserved capacity before calculating allocatable capacity.
Overcommit ratio	The multiplier used to calculate how much virtualized capacity can be allocated from physical resources. A higher overcommit ratio increases the logical resource or storage capacity that can be allocated to VMs.
Allocated capacity	The logical resource capacity that has been allocated to VMs or related resources.
Free to allocate	The remaining capacity after allocated capacity is deducted from the total allocatable capacity.

8.2.1 Configure Cluster Overcommit

Configure CPU overcommit ratio, memory overcommit ratio, and host reserved memory in the cluster advanced settings to control the computing resources available to VMs in a cluster.

Procedure

1. Go to the configuration entry of the target cluster.

   On the target cluster page, click **Actions** > **Advanced Settings** > **Edit Cluster Overcommit**. You can also go to the **Advanced Settings** tab of the cluster to modify the settings.

2. In the cluster advanced settings, configure computing resource overcommit options as needed.

Option	Description
CPU Overcommit Ratio	Controls the number of virtual CPUs that can be allocated to VMs. Allocatable virtual CPUs = physical CPU total threads in the cluster x CPU overcommit ratio.
Memory Overcommit Ratio	Controls the amount of virtual memory that can be allocated to VMs. Allocatable virtual memory capacity = (physical memory capacity - reserved capacity) x memory overcommit ratio.

3. Confirm the configuration and click **OK**.

8.2.2 Configure Storage Overcommit

Configure the data storage overcommit ratio in the data storage advanced settings to control the data storage space that can be allocated to VMs.

Context

The data storage overcommit ratio controls the allocatable data storage space for VMs. Local storage, NFS storage, SAN storage, and distributed storage all support storage overcommit configuration in their corresponding data storage settings.

Procedure

1. Go to the configuration entry of the target data storage.

   The configuration entries and configurable options vary by data storage type. For details about configuration entries for each data storage type, see:

   - *Modify Local Storage Configuration*
   - *Modify NFS Storage Configuration*
   - *Modify SAN Storage Configuration*
   - *Distributed Storage Configuration Management*

2. In the data storage advanced settings, set **Data Storage Overcommit Ratio**.

   The data storage overcommit ratio is used to calculate the allocatable capacity of data storage. The calculation formula is as follows:

   ```
   Data storage allocatable capacity = [(actual capacity - reserved capacity) x overcommit ratio] - (threshold capacity + sum of all disks allocated to VMs + snapshots + image cache + migration cache)
   ```

   Note:
   SAN storage also provides advanced settings such as disk preallocation policy and storage allocation strategy. When modifying SAN storage configuration, set the overcommit ratio based on the storage provisioning method and service capacity plan.

3. Confirm the configuration and click **OK**.

8.2.3 Resource Allocation and Overcommit Formulas

Capacity Calculation Formulas

Category	Formula
Resource utilization rate	Total CPU = physical cores x single-core GHz
Resource allocation ratio	• Allocation ratio = allocated : total overcommit capacity • Total overcommit capacity = physical total - reserved physical capacity • Total allocatable capacity = total overcommit capacity x overcommit ratio • Remaining allocatable capacity = total allocatable capacity - allocated capacity
CPU resource distribution	CPU overcommitted total = physical CPU total x overcommit ratio
Memory resource distribution	• Memory overcommitted total = reserved memory + total allocatable memory capacity • Total allocatable memory = (physical memory total - reserved memory) x overcommit ratio
Storage resource distribution	• Storage overcommitted total = reserved capacity + total allocatable storage capacity • Total allocatable storage = (physical storage total - reserved capacity) x overcommit ratio

Meaning of Overcommit and Allocation

- **CPU overcommit**: A single physical CPU core can be virtualized into up to N logical CPU cores for allocation to VMs. For example, if the CPU overcommit ratio is 2:1, one physical CPU core can be virtualized into up to two logical CPU cores. If a host has 10 physical CPU cores, up to 20 logical CPU cores can be allocated to VMs.
- **Memory/storage overcommit**: One unit of memory or storage capacity can be virtualized into up to N units of memory or storage capacity for allocation to VMs. For example, if the memory or storage overcommit ratio is 2:1, 1 GB of memory or storage capacity can be virtualized into up to 2 GB.

- **CPU allocation**: A physical CPU core is actually virtualized into N logical CPU cores for use by VMs. Therefore, the CPU allocation ratio is less than or equal to the CPU overcommit ratio. For example, if the CPU allocation ratio is 1.5:1, one physical CPU core is actually virtualized into 1.5 logical CPU cores.
- **Memory/storage allocation**: One unit of memory or storage capacity is actually virtualized into N units of memory or storage capacity. Therefore, the memory or storage allocation ratio is less than or equal to the memory or storage overcommit ratio.

Calculation Example

Using host storage as an example, if the total physical storage capacity is 100 GB, the reserved physical capacity is 10 GB, the overcommit ratio is 2:1, and the allocated capacity is 150 GB, then:

- Storage allocation ratio = 150 GB : 90 GB = 1.67
- Total overcommit storage = 100 GB - 10 GB = 90 GB
- Total allocatable storage = 90 GB x 2 = 180 GB
- Remaining allocatable storage = 180 GB - 150 GB = 30 GB

8.3 Cluster DRS

Dynamic Resource Scheduling (DRS) monitors CPU or memory load on hosts at the cluster level and dynamically adjusts VM services running on hosts according to the configured scheduling policy.

NexaSphere supports both manual and automatic scheduling policies. A manual scheduling policy provides scheduling suggestions, and administrators can manually migrate VMs according to the suggestions. An automatic scheduling policy allows the system to execute resource scheduling based on the scheduling algorithm. Both scheduling policies can balance cluster loads and improve platform stability.

8.3.1 Configure DRS

Prerequisites

- Only nSDS distributed storage or SAN storage data stores are attached to the cluster.
- All hosts in the cluster use the same CPU model.

Procedure

1. Go to the DRS configuration entry.
   - On the **DRS** tab of the target cluster, click **Enable**.

- In the navigation bar on the left side of the platform page, choose **Reliability** > **DRS Policy**. Then, for the target cluster, choose **Actions** > **Modify Policy**.

2. Configure the DRS policy.

Option	Description
DRS Mode	• Manual Scheduling: If the CPU or memory usage of the hosts in the cluster reaches the specified threshold, administrators manually schedule the resources based on the scheduling suggestions. • Automatic Scheduling: If the CPU or memory usage of the hosts in the cluster reaches the specified threshold, the system auto-schedules the resources based on the scheduling algorithm.
Monitoring Item	• CPU Utilization: Defines the trigger condition for the CPU utilization monitoring item. • Memory Utilization: Defines the trigger condition for the memory utilization monitoring item. • CPU or Memory Utilization: Triggers scheduling evaluation when either resource utilization reaches the trigger condition.
Duration	Defines how long the threshold must be met. The unit can be seconds, minutes, or hours. The maximum duration is 6 hours.
VM Migration Concurrency	Specifies the maximum number of VMs that can be migrated from their current hosts to recommended target hosts at the same time during DRS. The default value is 1, and the maximum value is 100.
Cluster Scanning Interval	Specifies the interval for scanning the cluster balance status during DRS. The default value is 10 minutes. The value cannot be less than 5 minutes or greater than 168 hours.

Note:

If any host in the cluster reaches the threshold of the monitoring item trigger condition and the configured duration is met, the cluster is considered imbalanced and scheduling suggestions are generated.

3. Confirm the configuration and click **OK**.

8.3.2 Perform DRS Operations

Context

After DRS is configured, you can scan the cluster balance status, execute scheduling suggestions , view scheduling tasks, reconfigure the policy, or disable DRS as needed.

- Manually scan the cluster balance status.
- When the scheduling policy is manual, migrate VMs to recommended hosts according to the scheduling suggestions to balance the cluster load.
- View scheduling tasks.

  Click **O&M Management** > **Task** > **Scheduling Task** to view scheduling execution history, execution results, execution time, and other information.
- Reconfigure the DRS policy.
- Disable Cluster DRS.

Note:
After DRS is disabled, host loads might not be balanced in high-load scenarios, which can affect service performance. Proceed with caution.

8.4 VM Scheduling Policy

VM scheduling policies define affinity or anti-affinity relationships between VMs, or between VMs and hosts. Administrators can use scheduling policies to control where VMs run and meet requirements such as service isolation, fault domain planning, performance optimization, and deployment on specific hosts.

Core Components

Scheduling policy	Defines the policy type and execution mechanism for VM scheduling.
VM Scheduling Group	The basic unit for associating VMs with scheduling policies.

- A VM can join only one VM scheduling group. After a VM joins a group, the VM is assigned to a host according to the scheduling policy associated with the group.
- A VM scheduling group can be associated with one or more scheduling policies.
- A scheduling policy can be associated with one VM scheduling group.
- Deleting a VM scheduling group also deletes the associated VM scheduling policies.

Host Scheduling Group

The basic unit for associating hosts with scheduling policies. Host scheduling groups are used only by VMs-exclusive-from-hosts and VMs-affinitive-to-hosts policies.

- A host can join only one host scheduling group. After a host joins a group, the host is scheduled according to the scheduling policy associated with the group.
- A host scheduling group can be associated with one or more scheduling policies.
- A scheduling policy can be associated with one host scheduling group.
- Deleting a host scheduling group also deletes the associated VM scheduling policies.

Policy Types and Execution Mechanisms

Type	Description
VM Exclusive from Each Other	VMs in the same VM scheduling group should not or must not run on the same host.
VM Affinitive to Each Other	VMs in the same VM scheduling group should or must run on the same host.
VMs Exclusive from Hosts	VMs in a VM scheduling group should not or must not run on hosts in a host scheduling group.
VMs Affinitive to Hosts	VMs in a VM scheduling group should or must run on hosts in a host scheduling group.

Each scheduling policy type supports the Hard and Soft execution mechanisms.

- The Hard mechanism requires VMs to be assigned to hosts strictly according to the policy. If no host satisfies the policy, VM startup fails.

- The Soft mechanism prioritizes host assignment according to the policy. If no host satisfies the policy, the VM can attempt to run on another host with sufficient resources.

How It Works

After a VM joins a VM scheduling group, the scheduling policy associated with the group affects host assignment for the VM.

- VM-exclusive-from-each-other and VM-affinitive-to-each-other policies do not require a host scheduling group. VMs are assigned to hosts according to the policy and its execution mechanism.
- VMs-exclusive-from-hosts and VMs-affinitive-to-hosts policies require both a VM scheduling group and a host scheduling group. VMs are assigned to hosts according to the policy and its execution mechanism.

Typical Scenarios

Different scheduling policies apply to different service isolation, resource aggregation, and host constraint scenarios.

- VM-exclusive-from-each-other policies are suitable for active/standby databases, distributed application nodes, and other scenarios where a single host failure must not affect multiple service instances.
- VM-affinitive-to-each-other policies are suitable for scenarios where multiple VMs are expected to run on the same host as much as possible to reduce communication latency.
- VMs-exclusive-from-hosts policies are suitable for scenarios where VMs need to avoid specified hosts, host groups, or fault domains.
- VMs-affinitive-to-hosts policies are suitable for scenarios where VMs need to be constrained to specified hardware specifications, rack locations, or host groups.

8.4.1 Create Mutually Exclusive/Affinitive VM Scheduling Policies

Prerequisites

Mutually exclusive or affinitive VM scheduling policies must be associated with a VM scheduling group before they can take effect on VMs in the group.

Procedure

1. In the navigation pane, choose **Reliability** > **VM Scheduling Policy** > **VM Scheduling Group**.
2. Click **New VM Scheduling Group** to create a VM scheduling group.

Option	Description
Name	Specifies the VM scheduling group name, for example, VM Scheduling Group A1.
Description	Enter a brief description for the VM scheduling group.
Virtual Machine	Selects one or more target VMs to join the scheduling group.

3. Click **OK**.
4. On the **VM Scheduling Policy** page, click **New VM Scheduling Policy**.
5. Create a mutually exclusive or affinitive VM scheduling policy.

Option	Description
Name	Specifies the VM scheduling policy name, for example, VM Scheduling Policy A1.
Description	Enter a brief description for the VM scheduling policy.
Type	Select **VM Exclusive from Each Other** or **VM Affinitive to Each Other**.
Execution Mechanism	• **Hard**: VMs are forcibly assigned to hosts based on the associated VM scheduling policies. If no host satisfies the policy, the VM fails to start. • **Soft**: VMs are preferentially assigned to hosts based on the associated VM scheduling policies. If no host satisfies the policy, the VM can attempt to run on another host with sufficient resources.
Associate VM Scheduling Group	After the policy is associated with a VM scheduling group, it takes effect on all VMs in the group. You can associate the policy with an existing scheduling group or create a scheduling group.

6. Click **OK**.

8.4.2 Create VMs Exclusive from/Affinitive to Hosts Scheduling Policies

Prerequisites

VMs exclusive from hosts or VMs affinitive to hosts scheduling policies must be associated with both a VM scheduling group and a host scheduling group before they can take effect.

Procedure

1. In the navigation pane, choose **Business Reliability** > **VM Scheduling Policy** > **VM Scheduling Group**.
2. Click **New VM Scheduling Group** to create a VM scheduling group.

Option	Description
Name	Specifies the VM scheduling group name, for example, VM Scheduling Group A2.
Description	Enter a brief description for the VM scheduling group.
Virtual Machine	Selects one or more target VMs to join the scheduling group.

3. On the **Host Scheduling Group** page, click **New Host Scheduling Group**.

Option	Description
Name	Specifies the host scheduling group name, for example, Host Scheduling Group A2.
Description	Enter a brief description for the host scheduling group.
Host	Selects one or more hosts to join the scheduling group.

4. On the **VM Scheduling Policy** page, click **New VM Scheduling Policy**.
5. Create a VMs exclusive from hosts or VMs affinitive to hosts scheduling policy.

Option	Description
Name	Specifies the VM scheduling policy name, for example, VM Scheduling Policy A2.
Description	Enter a brief description for the VM scheduling policy.
Type	Select **VMs Exclusive from Hosts** or **VMs Affinitive to Hosts**.
Execution Mechanism	• **Hard**: VMs are forcibly assigned to hosts based on the associated VM scheduling policies. If no host satisfies the policy, the VM fails to start.

Option	Description
	• **Soft**: VMs are preferentially assigned to hosts based on the associated VM scheduling policies. If no host satisfies the policy, the VM can attempt to run on another host with sufficient resources.
Associate VM Scheduling Group	After the policy is associated with a VM scheduling group, it takes effect on all VMs in the group.
Associate Host Scheduling Group	After the policy is associated with a host scheduling group, it applies the exclusive or affinitive rule based on the hosts in the group.

6. Click **OK**.

8.4.3 Manage VM Scheduling Policies and Related Resources

Context

After creating VM scheduling policies, VM scheduling groups, or host scheduling groups, you can adjust the enablement status, basic information, and group resources as needed.

- Enable or disable a VM scheduling policy.
  - If the scheduling policy is disabled, the scheduling state is **Inactive**. In this state, the scheduling policy does not take effect on the associated VMs.
  - If the scheduling policy is enabled, the scheduling state may be **Normal** or **Conflicted**. Normal indicates that all associated VMs are running on hosts assigned according to the scheduling policy. Conflicted indicates that some associated VMs are not running on hosts assigned according to the scheduling policy.
- Modify the name, description, or execution mechanism of a VM scheduling policy.
- Delete a VM scheduling policy that is no longer needed.

  Note:
  After a VM scheduling policy is deleted, the associated VMs are no longer scheduled according to that policy.
- Add VMs to or remove VMs from a VM scheduling group.
- Modify or delete a VM scheduling group.

Note:
After a VM scheduling group is deleted, the associated VM scheduling policies are also deleted. Proceed with caution.

- Add hosts to or remove hosts from a host scheduling group.
- Modify or delete a host scheduling group.

Note:
After a host scheduling group is deleted, the associated VM scheduling policies are also deleted. Proceed with caution.

8.5 VM HA Policy

HA Policy controls the automatic recovery behavior of VMs when a VM stops unexpectedly, is shut down as planned, or encounters related resource failures. Administrators can use the global HA policy, the VM HA switch, and failover policies to automatically restart critical services or migrate them to other hosts after a failure.

Core Concepts

VM HA	Specifies whether a VM automatically restarts after a planned or unexpected shutdown. If the platform HA policy is not enabled, turning on the VM HA switch takes effect after the platform HA policy is enabled.
Failover policy	Specifies whether a VM is migrated to another host and started when compute, storage, network, or other related resources of the VM fail.

Detection Items

- Management network connection status: Detects the network connection status between the host where the VM is located and the management node. A management node failure or management network interruption causes a management network connection status failure.
- Storage network connection status: Detects the network connection status between the VM and the data storage resource where its system disk is located. A failure of the data storage where the system disk is located or a storage network interruption causes a storage network connection status failure.
- Business NIC status: A failure of the host business NIC associated with the business VM or the switch port directly connected to the business NIC causes a business NIC failure for the VM.

Typical Failover Scenarios

Typical Scenario	Management Network Connection Status	Storage Network Connection Status	Business NIC Status	Migrate on Failure?
Scenario A	Normal	Normal	Failure	Migrate / Do Not Migrate
Scenario B	Normal	Failure	Normal	Migrate / Do Not Migrate
Scenario C	Normal	Failure	Failure	Migrate / Do Not Migrate
Scenario D	Failure	Normal	Normal	Do Not Migrate

How It Works

When a VM stops because of an unexpected condition or a planned shutdown, NexaSphere checks the VM HA switch status. If the switch is turned on, the system restarts the VM on the current host or another host.

When the management network connection status, storage network connection status, or business NIC status of the host where a VM is located becomes abnormal, NexaSphere checks the fault migration policy and the VM HA switch status. If the corresponding fault migration switch and the VM HA switch are both turned on, the VM is migrated to another host and started.

8.5.1 Configure HA Policy

Procedure

1. In the navigation pane, choose **Reliability** > **HA Policy**.
2. Enable HA Policy.

   HA Policy is enabled by default. If it has been disabled, turn on the switch at the top of the **HA Policy** page.
3. On the **Migration Policy** page, set the failover policy.

Typical Scenario	Management Network Connection Status	Storage Network Connection Status	Business NIC Status	Migrate on Failure?	Migration Description
Scenario A	Normal	Normal	Failure	Migrate / Do Not Migrate	Supports setting the policy to migrate or do not migrate.
Scenario B	Normal	Failure	Normal	Migrate / Do Not Migrate	Supports setting the policy to migrate or do not migrate. In a SAN storage environment, if this policy is set to do not migrate, a storage network connection status failure still triggers automatic migration.
Scenario C	Normal	Failure	Failure	Migrate / Do Not Migrate	The migration policy for simultaneous storage connection status and business NIC status failures follows the migration policy configured for either individual failure: • If the migration policy for both storage connection status failure and business NIC status failure is do not migrate, this policy is set to do not migrate. • If the migration policy for either failure scenario is migrate, this policy is set to migrate.
Scenario D	Failure	Normal	Normal	Do Not Migrate	When the management network status is faulty, setting a fault migration policy is not supported.

Note:

Storage network connection status detection supports only shared storage and does not currently support local storage.

4. On the **Migration Policy** page, set the host error detection policy.

Host Fault Detection Item	Description
Host Self-Inspection Interval	Specifies the interval for host status self-inspection. Default: 5 seconds.
Maximum Host Self-Inspection Attempts	Specifies the maximum number of host self-inspection attempts. If all attempts fail, the system determines that a network error has occurred on the host. Default: 6 attempts.

5. On the **Advanced Settings** page, configure advanced HA policy settings as needed.

Category	Name	Description
Virtual Machine	Maximum Interval for VM HA Restart Attempts	Specifies the maximum interval for the system to execute the GC task and retry starting a VM after the VM stops unexpectedly. Default: 300 seconds.
	VM HA Restart Retry Delay	Specifies the delay before the next restart attempt if the previous restart attempt fails for a VM with HA enabled. Default: 60 seconds.
	HA VM State Scanning Interval	Specifies the interval for scanning the state of a VM with HA enabled after restart fails. Default: 60 seconds.
	HA VM State Update Speed	• Specifies the response speed for updating the UI state list when the enabled state of an HA VM changes. Default: 1. Valid values: -1 to 5. • A smaller value indicates faster response . A larger value indicates slower response , but the system ignores more outdated state change notifications and reduces system load. • -1 indicates that the UI state list is not actively updated.
Host	Timeout for Host Connection to Data Storage	Specifies the timeout for a host to check its connection to data storage. Default: 5 seconds.

Category	Name	Description
	Abnormal Host Status Update Interval	Specifies the interval for checking and updating the status of abnormal hosts. Default : 5 seconds.
	Connection Attempts to Determine Host Disconnect ion	Specifies the maximum number of attempts to reconnect to a host. If all attempts fail, the system determines that the host is disconnect ed. Default: 12 attempts.
	Ping Response Time to Determine Successful Host Connection	Specifies the time for determining whether a host is successfully connected to the management node by ping. If the host responds within this time, the connection is considered successful. Default: 5 seconds.
	Connection Success Rate to Determine Host Recovery	Specifies the success rate used to determine whether the host has reconnected to the management node when both successful and failed connection attempts occur within the specified number of attempts. Default: 50%.
	Minimum Successful Connections to Determine Host Recovery	Specifies the minimum number of successful connections that a host must establish with the management node before the host is determined to be reconnected. Default: 5.

8.5.2 View HA Tasks

Context

After HA Policy is enabled, HA task records are generated when the platform triggers the HA mechanism.

Procedure

1. In the navigation pane, choose **O&M Management** > **Task** > **HA Task**.
2. View HA task information.

   The HA task page shows task result, VM name, VM owner, previous host, target host, start time, and completion time, which can be used to audit and trace HA execution.

8.5.3 Disable HA Policy

Procedure

1. In the navigation pane, choose **Reliability** > **HA Policy**.

2. On the **HA Policy** page, click **Disable**.

> **Note:**
> After HA Policy is disabled, VMs no longer automatically restart after shutdown, which might cause service interruptions. Proceed with caution.

8.5.4 HA Policy Practice

Context

Assume that four business VMs are deployed on Host A to run MySQL database services. To ensure service high availability, all four VMs must be migrated to another host when the business NIC of Host A fails. In this scenario, set the HA mode of these VMs to **NeverStop**, configure migration to be triggered when the business NIC status fails, and ensure that sufficient host resources are available in the platform.

Procedure

1. Enable HA Policy.

   In the navigation pane, choose **Reliability** > **HA Policy**. Turn on the switch at the top of the **HA Policy** page.

2. Turn on the HA switch for VMs.

   The VM HA switch can be configured at the VM level or cluster level. The VM-level configurat ion takes precedence over the cluster-level configuration.

   - When creating a VM, turn on the **HA** switch.
   - Go to the cluster where the VM resides, choose **Modify Configuration** > **Advanced Settings** > **VM Settings**, and turn on the **VM HA** switch. After this setting is enabled, the VM HA switch is turned on by default for new VMs created in the cluster.

3. Configure the VM fault migration policy.

   Go to the **HA Policy** > **Migration Policy** page and turn on the **Failover** switch for Scenario A. When this switch is turned on, the **Failover** switch for Scenario C is automatically turned on.

Result

After the VM HA policy is configured, if the business NIC of Host A fails, the four VMs on Host A are automatically migrated to Host B and started. You can search for related task records in **O&M Management** > **Task** > **HA Task**.

9 Monitoring and Operations

9.1 Overview

NexaSphere provides centralized operations capabilities for platform status, resource performance, capacity, health, alarms, tasks, events, and logs. Operations administrators can use the dashboard and resource detail pages to understand overall platform status, use monitoring charts and capacity information to identify resource pressure, use alarms and notifications to track exceptions, and use tasks, events, and logs to trace platform operations and troubleshooting clues.

SNMP monitoring and notification integration connects monitoring data and alarm messages from NexaSphere to third-party systems. Tags and custom attributes help mark, filter, and search resources during routine inspection, troubleshooting, and resource governance.

9.2 Resource Overview and Capacity Monitoring

NexaSphere provides dashboard resource statistics, data center resource topology, and resource capacity information to help administrators quickly understand overall platform status, resource distribution, and capacity pressure. Administrators can first use the dashboard to view resource status and alarm statistics for a specified data center, and then open resource overview pages to view resource hierarchy and capacity usage.

9.2.1 View Dashboard Resource and Alarm Statistics

View Dashboard Resource Statistics by Data Center

On the **Dashboard** page, select a data center to view core physical resources, virtual resources, recent access information, and alarm statistics for the data center.

Refresh and Display Rules

- Each time you open or refresh the dashboard, the page fetches and displays the latest data.
- Chart modules automatically refresh data every 30 seconds by default.
- Status charts use a consistent color scheme: green indicates normal status, red indicates abnormal status, and gray indicates other statuses.
- Percentage progress bars use blue, yellow, and red to indicate different usage levels. Blue indicates less than 60%, yellow indicates greater than or equal to 60% and less than 80%, and red indicates greater than or equal to 80%.

- Resource status cards and some utilization trend and usage statistics cards support navigation to the corresponding resource pages.

Resource Usage Statistics

The dashboard displays usage statistics for VMs, hosts, data storage, and image storage by resource type. Some cards support navigation to the corresponding resource page by clicking the resource name or statistics.

Card	Statistics
Virtual Machine	Displays the total number of VMs, running VMs, stopped VMs, and VMs in other ready statuses.
Top 10 Virtual Machine CPU Utilization	Displays the top 10 VMs in the data center by CPU utilization. A VM must have VMTools installed to be included in the ranking.
Top 10 Virtual Machine Memory Utilization	Displays the top 10 VMs in the data center by memory utilization. A VM must have VMTools installed to be included in the ranking.
Host	Displays the total number of hosts, connected hosts, disconnected hosts, and hosts in other ready statuses.
CPU Utilization	Displays the total number of logical CPU cores and the average CPU utilization across all available hosts.
Memory Utilization	Displays the average memory utilization, used memory, and total memory across all available hosts.
Top 3 Host CPU Utilization	Displays the top 3 hosts in the data center by CPU utilization.
Top 3 Host Memory Utilization	Displays the top 3 hosts in the data center by memory utilization.
Data Storage Utilization	Displays data storage utilization, remaining available capacity, used capacity, and total capacity. You can view information for all data storage or a single data storage resource.
Image Storage Utilization	Displays image storage utilization, remaining available capacity, used capacity, and total capacity. You can view information for all image storage or a single image storage resource.

Alarm Statistics

The dashboard displays unread alarm statistics for the past seven days, helping administrators quickly identify exceptions that require attention in the current data center.

Dimension	Description
Alarm level	Displays the number of unread alarms in the past seven days by Emergent, Major, and Info levels.
Alarm event	Displays the event name, occurrence time, alarm count, and alarm level for events that have triggered alarms. Alarm levels are color-coded, and you can click an event to view alarm details.

9.2.2 View Data Center Resource Topology

The data center resource topology shows the resource hierarchy and summary information in a data center.

Procedure

1. In the navigation pane, choose **Inventoy** > **VM and Host**. Then, select the target data center.
2. On the data center details page, open the **Overview** tab.
3. View the data center resource topology.

Result

The resource topology shows the organization of clusters, hosts, networks, data storage, VMs , and other resources in the data center, helping administrators quickly understand resource distribution.

9.2.3 View Resource Capacity Information

NexaSphere displays capacity usage and allocation information for compute and storage resources in the **Capacity Info** card on resource overview pages. Administrators can view resource capacity status from root node, data center, cluster, host, data storage, and VM levels.

Object	Monitoring Metrics and Description
Root Node	Displays the total physical CPU GHz and average utilization, total physical memory and average utilization, total physical storage and average utilization, and remaining available memory and storage across all data centers.

Object	Monitoring Metrics and Description
Data Center	Displays the total physical CPU GHz and average utilization, total physical memory and average utilization, total physical storage and average utilization, and remaining available memory and storage in the data center.
Data Storage	Displays storage utilization, storage allocation ratio, and storage distribution, including total storage resources, utilization, remaining available capacity, total capacity after overcommit, reserved capacity, allocated capacity, and remaining allocatable capacity.
Cluster	Displays utilization, allocation ratio, and resource distribution for all physical CPU and physical memory resources in the cluster.
Host	Displays utilization, allocation ratio, and resource distribution for physical CPU, physical memory, and physical storage resources on the host.
VM	Displays the number of CPU cores and CPU utilization, total memory, used memory, remaining available memory, total storage, used storage , and remaining available storage for the VM.

9.3 Resource Performance Monitoring

NexaSphere provides visual charts that display resource performance monitoring data over a period of time, helping administrators understand compute, storage, and network utilization trends and evaluate resource health.

Monitoring Chart Types

Chart Type	Description
Bar Chart	Displays capacity usage data as proportional bars, helping administrators quickly understand resource capacity information.
Line Chart	Displays resource performance metrics as line charts, helping administrators observe metric trends over time.

Monitoring Data Collection Methods

NexaSphere provides Basic Monitoring and Advanced Monitoring for VMs. Advanced Monitoring usually provides more accurate memory data.

Basic Monitoring Monitoring data is obtained from the host through Libvirt.

Advanced Monitoring	Monitoring data is obtained from inside the VM by the Advanced Monitoring Agent. VMTools must be installed in the VM before Advanced Monitoring can be used.

Monitoring Data Collection Intervals

NexaSphere uses real-time monitoring. Resource monitoring charts refresh data every 10 seconds by default.

9.3.1 View Monitoring Charts

NexaSphere supports viewing resource utilization metrics through line charts.

Procedure

1. In the navigation pane, choose **Inventory**.
2. Select the resource object to view, such as a VM, host, cluster, image storage, data storage, or distributed port group.
3. On the resource details page, click **Monitoring**.
4. Select the metrics to display as needed.
5. Select or customize the time range as needed.
6. Select one or more monitoring objects as needed.

9.3.2 Customize Monitoring Charts

Administrators can adjust the time range, metrics, monitoring objects, and chart position displayed in monitoring charts.

- View data point details.

  Hover over a chart to view details about the corresponding data point.
- Adjust the time range.

  By default, monitoring data from the last 15 minutes is displayed. Valid values include 15 minutes, 1 hour, 6 hours, 1 day, 1 week, 1 month, 1 year, and a custom time range.
- Select metrics.

  Select the metrics to focus on based on service requirements.
- Select monitoring objects.

  Display data for all monitoring objects or for specified monitoring objects.
- Adjust chart positions.

  Drag monitoring charts to change their display order.

9.3.3 Metrics Reference

Supported metrics vary by resource object. The following table lists common resource objects and metrics.

Object	Metric	Item and Description
Cluster	CPU, memory, disk, NIC	CPU utilization sum, memory usage percentage, disk IOPS sum (read/write), and NIC data transfer rate sum (send/receive).
Host	CPU	• **CPU Utilization**: The proportion of time the CPU is in a non-idle state. • **CPU Idle Rate**: The proportion of time the CPU is in an idle state. • **CPU Occupancy Rate (System Process)**: The proportion of time the CPU spends in kernel space for operations such as memory allocation, I/O operations, and child process creation. • **CPU Occupancy Rate (User Process)**: The proportion of time the CPU spends in user space for programs such as shells, databases, and web servers. • **CPU Occupancy Rate Average (Waiting)**: The proportion of time the CPU spends waiting for the hard disk drive to load data into memory after initiating a read or write operation.
Host	Memory	**Memory Usage**: The amount of used and free resource memory.
Host	Disk	• **Disk Speed**: The read and write speed of the resource disk. • **Disk IOPS**: The read and write IOPS of the resource disk. • **Disk Latency**: The latency of the resource disk. • **Total Disk Usage Ratio**: The percentage of used capacity across all host disks. • **Total Disk Usage**: The amount of used capacity across all host disks. • **Disk Usage Ratio of Platform System Files**: The percentage of disk capacity occupied by platform system files.

Object	Metric	Item and Description
		• **Disk Usage of Platform System Files**: The amount of disk capacity occupied by platform system files.
Host	LUN	• **LUN Speed: Read/Write**: The current read and write throughput of the LUN. • **LUN IOPS: Read/Write**: The number of read and write I/O operations completed by the LUN per second. • **LUN Latency**: The I/O response latency of the LUN.
Host	NIC	• **NIC Data Transfer Rate**: The current send and receive rate of the resource NIC. • **NIC Packet Rate**: The current send and receive packet rate of the resource NIC. • **NIC Packet Discard Rate**: The current packet drop rate for outgoing and incoming packets on the resource NIC.
VM	CPU	• **CPU Utilization**: The proportion of time the CPU is in a non-idle state. • **CPU Idle Rate**: The proportion of time the CPU is in an idle state. • **CPU Occupancy Rate (System Process)**: The proportion of time the CPU spends in kernel space. • **CPU Occupancy Rate (User Process)**: The proportion of time the CPU spends in user space. • **CPU Occupancy Rate Average (Waiting)**: The proportion of time the CPU spends waiting for disk I/O.
VM	Memory	• **Memory Usage**: The amount of used and free resource memory. • **Available Memory Capacity**: The available amount of resource memory that can be used. • **Free Memory Capacity**: The amount of free resource memory. • **Total Memory Capacity**: The total amount of resource memory. • **Memory Idle Rate**: The percentage of resource memory currently in an idle state.

Object	Metric	Item and Description
		• **Memory Utilization**: The percentage of resource memory currently in use.
VM	Disk	• **Disk Speed**: The read and write speed of the resource disk. • **Disk IOPS**: The read and write IOPS of the resource disk. • **Disk Utilization**: The percentage of used capacity on the resource disk. • **Disk Idle Rate**: The percentage of idle capacity on the resource disk. • **Disk Usage Capacity**: The amount of used capacity on the resource disk. • **Disk Idle Capacity**: The amount of free capacity on the resource disk.
VM	NIC	• **NIC Data Transfer Rate**: The current send and receive rate of the resource NIC. • **NIC Packet Rate**: The current send and receive packet rate of the resource NIC. • **NIC Packet Discard Rate**: The current packet drop rate for outgoing and incoming packets on the resource NIC.
Data Storage	Capacity	**Capacity Percent Used**: The percentage of capacity currently used by the resource.
Standalone Image Storage/ Distributed Image Storage	Capacity	**Capacity Percent Used**: The percentage of capacity currently used by the resource.
Standalone Image Storage	CPU, disk, memory, NIC	• **CPU Utilization**: The proportion of time the CPU is in a non-idle state. • **CPU Idle Rate**: The proportion of time the CPU is in an idle state. • **CPU Occupancy Rate (System Process)**: The proportion of time the CPU spends in kernel space, performing typical operations such as memory allocation, I/O operations, and creating child processes. • **CPU Occupancy Rate (User Process)**: The proportion of time the CPU spends in user space,

Object	Metric	Item and Description
		running typical user-space programs such as shells, databases, and web servers. • **CPU Occupancy Rate Average (Waiting)**: The proportion of time the CPU spends waiting for the hard disk drive to load data into memory after initiating a read or write operation. • **Disk Speed**: The read and write speed of the resource disk. • **Disk IOPS**: The read and write IOPS of the resource disk. • **Memory Usage**: The amount of used and free resource memory. • **NIC Data Transfer Rate**: The current send and receive rate of the resource NIC. • **NIC Packet Rate**: The current send and receive packet rate of the resource NIC. • **NIC Packet Discard Rate**: The current packet drop rate for outgoing and incoming packets on the resource NIC.
Distributed Port Group	IP	• **Used IP Percentage (IPv4)**: The percentage of IPv4 addresses currently used by the resource. • **Available IP Percentage (IPv4)**: The percentage of remaining available IPv4 addresses on the resource.

9.4 Platform and Hardware Health Monitoring

9.4.1 View Management Node Monitoring

View management node monitoring data on the **Reliablility** > **MN Monitoring** page.

Status Description

- The page uses green, red, and gray to indicate status. Green indicates normal status. Red and gray indicate abnormal status.
- Management node HA works in active/standby mode. The node that displays the VIP is the active management node, and the node without a VIP is the standby management node.
- If all standby management nodes are abnormal, failover cannot occur when the active management node fails, which may cause management node downtime.

Management Service Metrics

Arbiter Gateway Reachable	Monitors whether the arbitration IP addresses of the active and standby management nodes are reachable. If the arbitration IP address is unreachable, management node HA may become unavailable.
Peer MN Reachable	Monitors whether the standby management node is reachable. If the standby management node is unreachable, the active management node cannot communicate with the standby management node.
VIP Reachable	Monitors whether the VIP is reachable. If the VIP is unreachable, the active management node cannot access the UI through the VIP.
Database status	Monitors database status. If the database is abnormal, data loss may occur and the fault must be recovered promptly.

9.4.2 View Host Hardware Monitoring

NexaSphere can monitor host hardware component status to help administrators identify host hardware exceptions.

Supported Hardware

- CPU
- Memory
- Physical disk
- Physical NIC
- GPU device
- Block device
- USB device
- Sensors, including voltage, current, fan, and temperature sensors.
- Power supply
- PCIe device

9.5 Alarms and Notifications

NexaSphere supports alarms for resource utilization, capacity usage, and predefined platform events. When core resources become abnormal, the platform can push alarm messages to specified endpoints so that administrators can identify and handle issues in time.

Alarm Service Infrastructure

Monitoring System	Provides time-series monitoring and alarms for resource utilization and capacity usage, collects predefined platform events and triggers event alarms, and supports custom alarms and alarm message templates.
Notification Service	Pushes alarm messages to system endpoints, email, DingTalk, Lark, WeCom, HTTP applications, Microsoft Teams, SNMP Trap receivers, and other endpoints.

Configuration Path

1. Prepare management node resources and notification channels required by the Alarm Service.
2. Create endpoints to receive alarm messages.
3. Create or adjust message templates to standardize alarm message formats.
4. Create resource alarms or event alarms and bind endpoints.
5. View, acknowledge, silence, or restore alarm messages.

9.5.1 Prepare the Alarm Service

The Alarm Service depends on monitoring data collection, storage, and notification channels. Before enabling a large number of monitoring and alarm configurations, plan management node resources and alarm message delivery channels in advance.

Resource Planning Recommendations

- Plan a dedicated physical server as the management node for the platform.
- Because monitoring data may periodically consume system disk I/O resources, use an SSD as the system disk of the management node.
- To avoid high system disk utilization caused by large amounts of monitoring data, plan more than 1 TB of system disk space.
- If the system disk space is less than 500 GB, shorten the monitoring data retention period and set an appropriate monitoring data retention size in *system parameters*.

Notification Channel Preparation

Before creating alarms, prepare external message receivers based on the required notification method, such as *email servers*, group chat robots, HTTP applications, SMS services, Microsoft Teams, or SNMP Trap receivers.

9.5.2 Endpoints

9.5.2.1 Create an Email Endpoint

Prerequisites

An email server has been added. For more information, see *Add an Email Server*.

Procedure

1. In the navigation pane, choose **O&M Management** > **Endpoint**.
2. On the **Endpoint** page, click **New Endpoint**.
3. Configure the email endpoint.

Option	Description
Name	Set the endpoint name.
Description	Enter a brief description for the endpoint.
Type	Select **Email**.
Email Server	Select an email server that has been added.
Email Address	Enter the email address that receives alarm messages.
Message Language	Set the notification language for alarm messages. Simplified Chinese and English are supported.

4. Review the configuration and click **OK**.

What's next

After creating an endpoint, you can continue to configure message templates and alarms.

9.5.2.2 Create a DingTalk Endpoint

Prerequisites

- An Advanced license has been installed and is valid.
- A DingTalk group bot has been added and its webhook URL has been obtained.

Procedure

1. In the navigation pane, choose **O&M Management** > **Endpoint**.
2. On the **Endpoint** page, click **New Endpoint**.
3. Configure the DingTalk endpoint.

Option	Description
Name	Set the endpoint name.
Description	Enter a brief description for the endpoint.
Type	Select **DingTalk**.
Address	Enter the webhook URL of the DingTalk bot.
Security Settings	Select the security settings configured for the DingTalk bot, including Signature or Other (keywords or IP addresses). • Custom Keywords: Alarm messages must contain at least one custom keyword to be sent successfully. If you choose this method, make sure you add "Alarm" as the keyword. Otherwise, alarm messages will fail to send. • IP Address: Only requests from within the specified IP address range will be processed by third-party applications. If you choose this method, add the management node IP address and VIP of the platform to the bot's IP allowlist to ensure that third-party applications receive alert messages correctly.
Mention Member	Specify whether to mention group members after the alarm message is pushed. Options include Nobody, All, or Specified Members. When choosing Specified Members, add the mobile phone numbers of the group members.
Message Language	Set the notification language for alarm messages. Simplified Chinese and English are supported.

4. Review the configuration and click **OK**.

9.5.2.3 Create a Lark Endpoint

Prerequisites

- An Advanced license has been installed and is valid.
- A Lark group bot has been added and its webhook URL has been obtained.

Procedure

1. In the navigation pane, choose **O&M Management** > **Endpoint**.
2. On the **Endpoint** page, click **New Endpoint**.
3. Configure the Lark endpoint.

Option	Description
Name	Set the endpoint name.
Description	Enter a description for the endpoint.
Type	Select **Lark**.
URL	Enter the webhook URL of the Lark bot.
Security Settings	Select the security settings configured for the Lark bot, including Signature or Other (keywords or IP addresses). • Custom Keywords: Alarm messages must contain at least one custom keyword to be sent successfully. If you choose this method, make sure you add "Alarm" as the keyword. Otherwise, alarm messages will fail to send. • IP Address: Only requests from within the specified IP address range will be processed by third-party applications. If you choose this method, add the management node IP address and VIP of the platform to the bot's IP allowlist to ensure that third-party applications receive alert messages correctly.
Mention Member	Specify whether to mention members when sending alarm messages through the bot. Options include Nobody, All, or Specified Members. When choosing Specified Members, add the user IDs of the designated members.
Message Language	Set the notification language for alarm messages. Simplified Chinese and English are supported.

4. Review the configuration and click **OK**.

9.5.2.4 Create a WeCom Endpoint

Prerequisites

- An Advanced license has been installed and is valid.
- A WeCom group bot has been added and its webhook URL has been obtained.

Procedure

1. In the navigation pane, choose **O&M Management** > **Endpoint**.
2. On the **Endpoint** page, click **New Endpoint**.
3. Configure the WeCom endpoint.

Option	Description
Name	Set the endpoint name.
Description	Enter a brief description for the endpoint.
Type	Select **WeCom**.
URL	Enter the webhook URL of the WeCom bot.
Mention Member	Specify whether to mention members when sending alarm messages through the bot. Options include Nobody, All, or Specified Members. When choosing Specified Members, add the user IDs of the designated members.
Message Language	Set the notification language for alarm messages. Simplified Chinese and English are supported.

4. Review the configuration and click **OK**.

9.5.2.5 Create an SMS Endpoint

Prerequisites

An AccessKey that includes the SMS service has been requested from the third-party platform.

Procedure

1. In the navigation pane, choose **O&M Management** > **Endpoint**.
2. On the **Endpoint** page, click **New Endpoint**.
3. Configure the SMS endpoint.

Option	Description
Name	Set the endpoint name.
Description	Enter a brief desctiption for the endpoint.

Option	Description
Type	Select **SMS**.
AccessKey ID	Enter the AccessKey ID that identifies the user.
AccessKey Secret	Enter the secret key used to authenticate the user.
SMS Address	Enter the mobile phone number that receives SMS messages.

4. Review the configuration and click **OK**.

9.5.2.6 Create an HTTP Application Endpoint

Prerequisites

The webhook URL of the HTTP application has been prepared, and the username and password have been configured as needed.

Procedure

1. In the navigation pane, choose **O&M Management** > **Endpoint**.
2. On the **Endpoint** page, click **New Endpoint**.
3. Configure the HTTP application endpoint.

Option	Description
Name	Set the endpoint name.
Description	Enter a brief description for the endpoint.
Type	Select **HTTP Application**.
URL	Enter the HTTP service URL.
Username	Enter the username configured for the HTTP application.
Password	Enter the password for the username.

4. Review the configuration and click **OK**.

9.5.2.7 Create a Microsoft Teams Endpoint

Prerequisites

- An Advanced license has been installed and is valid.
- The Incoming Webhook app has been added in Microsoft Teams and its webhook URL has been obtained.

Procedure

1. In the navigation pane, choose **O&M Management** > **Endpoint**.
2. On the **Endpoint** page, click **New Endpoint**.
3. Configure the Microsoft Teams endpoint.

Option	Description
Name	Set the endpoint name.
Description	Enter a brief description for the endpoint.
Type	Select **Microsoft Teams**.
URL	Enter the webhook URL obtained from Microsoft Teams.
Message Language	Set the notification language for alarm messages. Simplified Chinese and English are supported.

4. Review the configuration and click **OK**.

9.5.2.8 Create an SNMP Trap Receiver Endpoint

Prerequisites

- An Advanced license has been installed and is valid.
- SNMP Management has been enabled and an SNMP Trap receiver has been added.

Procedure

1. In the navigation pane, choose **O&M Management** > **Endpoint**.
2. On the **Endpoint** page, click **New Endpoint**.
3. Configure the SNMP Trap receiver endpoint.

Option	Description
Name	Set the endpoint name.
Description	Enter a brief description for the endpoint.
Type	Select **SNMP Trap Receiver**.
SNMP Trap Receiver	Select an SNMP Trap receiver that has been added.

4. Review the configuration and click **OK**.

9.5.2.9 Manage Endpoints

Context

After creating an endpoint, you can modify the endpoint, control its enabled state, adjust associated alarms, or delete the endpoint as needed.

- Modify endpoint basic information or configuration.

  On the **Endpoint** page, click **Actions** > **Edit Name and Description** to modify the name and description, or click **Actions** > **Modify Configuration** to modify endpoint configuration.
- Enable or disable endpoints.

  On the **Endpoint** page, select an endpoint and click **Enable** or **Disable**.
- Add or remove alarms for an endpoint.

  On the **Endpoint** page, select an endpoint, click **Actions** > **Add Alarm/Remove Alarm**, and then select the alarms to add or remove.
- Delete an endpoint.

  On the **Endpoint** page, select an endpoint and click **Actions** > **Delete**.

  Note:
  The system endpoint cannot be deleted.

9.5.3 Message Templates

9.5.3.1 Create a Message Template

A message template defines the text format used when an alarm pushes alarm messages to endpoints.

Prerequisites

Before creating DingTalk, Lark, WeCom, or Microsoft Teams message templates, make sure an Advanced license has been installed and is valid.

Procedure

1. In the navigation pane, choose **O&M Management** > **Message Template**.
2. On the **Message Template** page, click **New Message Template**.
3. Configure basic information for the message template.

Option	Description
Name	Set the message template name.

Option	Description
Description	Enter a brief description for the message template.

4. Configure template information.

   Email, DingTalk, Lark, WeCom, HTTP application, and Microsoft Teams templates support the following options.

Option	Description
Type	Select Email, DingTalk, Lark, WeCom, HTTP Application, or Microsoft Teams.
Alarm Type	Select resource alarm or event alarm.
Alarm Message Title	Set the title displayed in alarm messages.
Alarm Message Text	Set the text displayed in alarm messages.
Recovery Message Title	Set the recovery message title sent when a resource recovers from the alarm state. This option is supported only for resource alarms.
Recovery Message Text	Set the text displayed in recovery messages. This option is supported only for resource alarms.
Default Template	Select this option to set the current message template as the default template.

   SMS templates support the following options.

Option	Description
Signature Name	Enter the SMS signature name requested from the third-party platform.
Resource Alarm Template	Set the resource alarm message template and enter the resource alarm template CODE.
Event Alarm Template	Set the event alarm message template and enter the event alarm template CODE.

Option	Description
Default Template	Set this template as the default template. After this is configured, all SMS messages are sent in this template format.

Note:

- Email and Lark alarm message templates must comply with Text syntax requirements.
- DingTalk and WeCom alarm message templates must comply with Markdown syntax requirements.
- HTTP application alarm message templates must comply with JSON syntax requirements.
- Microsoft Teams alarm message templates must comply with Microsoft Teams official webhook syntax requirements.
- For SMS alarm message templates, third-party SMS signatures and SMS templates must be requested in advance. Currently, only Alibaba Cloud SMS Service integration is supported. Any template changes require re-applying through the third-party service.

5. Review the configuration and click **OK**.

9.5.3.2 Manage Message Templates

Context

After creating a message template, you can modify its basic information, set or unset the default template, or delete the template as needed.

- Modify basic information.

  On the **Message Template** page, select a message template and click **Actions** > **Edit Name and Description**.

- Set or unset the default template.

  On the **Message Template** page, select a message template and click **Set as Default** or **Unset Default**.

- Delete a message template.

  On the **Message Template** page, select a message template and click **Delete**.

9.5.4 Alarms

9.5.4.1 Alarm Rules

Resource Alarm Rules

Option	Description	Example
Resource Type	The resource type monitored by the alarm.	VM
Metric Item	The monitoring metric type and name.	CPU Utilization
Resource	The resource object monitored by the alarm.	/
Operator	The comparison relationship between the detected metric value and the threshold, including >, ≥, <, and ≤.	>
Threshold	The threshold and unit for triggering the alarm.	85%
Duration	Sends an alarm message after the alarm is continuously triggered for the specified time. Valid values include 30 seconds, 1 minute, 5 minutes, 10 minutes, 30 minutes, 1 hour, and custom.	5 minutes
Alarm Interval	The interval for repeated notifications after an alarm is generated, including Only Once, Every 1 Hour, and custom.	Every 1 Hour
Severity	The alarm message level, including Emergent, Major, and Minor.	Major
Recovery Notification	Sends one recovery notification when the monitored resource recovers from the alarm state to normal state.	/

Event Alarm Rules

Option	Description	Example
Resource Type	The resource type monitored by the alarm.	Data Storage
Metric Item	The monitoring metric name.	Data Storage Disconnected

Option	Description	Example
Severity	The alarm message level, including Emergent, Major, and Minor.	Emergent

9.5.4.2 Create a Resource Alarm

Resource alarms monitor time-series data of resources in the platform.

Prerequisites

- You can use *system parameters* to control default alarm behavior.
- Some alarm items require VMTools to be installed on VMs before monitoring and alarming can be performed.

Procedure

1. In the navigation pane, choose **O&M Management** > **Alarm** > **Resource Alarm**.
2. On the **Resource Alarm** page, click **New Resource Alarm**.
3. Configure the resource alarm.

Option	Description
Name	Set the resource alarm name.
Description	Enter a brief description for the resource alarm.
Resource Type	Select a resource type that supports alarms.
Metric Item	Select an alarm item based on the resource type. Some alarm items require you to select resources and set alarm trigger rules.
Alarm Interval	Select an alarm interval type, including Only Once and Repeat.
Severity	Set the alarm message level, including Emergent, Major, and Minor.
Alarm Recovery Notification	After this option is enabled, the system sends one recovery notification when any resource monitored by the alarm recovers from the alarm state.
Endpoint	Select the endpoint that receives alarm messages. A system alarm endpoint is provided by default, and custom endpoints can also be selected.

4. Review the configuration and click **OK**.

9.5.4.3 Create an Event Alarm

Event alarms monitor predefined events in the platform.

Procedure

1. In the navigation pane, choose **O&M Management** > **Alarm** > **Event Alarm**.
2. On the **Event Alarm** page, click **New Event Alarm**.
3. Configure the event alarm.

Option	Description
Resource Type	Select a resource type that supports event alarms.
Metric Item	Select an event alarm item based on the resource type.
Severity	Set the alarm message level, including Emergent, Major, and Minor.
Endpoint	Select the endpoint that receives alarm messages.

4. Review the configuration and click **OK**.

9.5.4.4 Manage Alarms

Context

After creating an alarm, you can modify, enable, disable, adjust endpoints, or delete the alarm as needed.

- Modify an alarm.

  On the **Alarm** page, click **Actions** > **Edit Name and Description** to modify the name and description, or click **Actions** > **Modify Configuration** to modify alarm configuration.
- Enable or disable alarms.

  On the **Alarm** page, select an alarm and click **Enable** or **Disable**.
- Add or remove endpoints for an alarm.

  On the **Alarm** page, click **Actions** > **Add Endpoint/Remove Endpoint**, and then select the endpoints to add or remove.
- Delete an alarm.

  On the **Alarm** page, select an alarm and click **Delete**.



Note:

Default alarms cannot be deleted.

9.5.5 Alarm Messages

9.5.5.1 View Alarm Messages

Context

Triggered alarm messages can be viewed on the unified Alarm Message page, on specific resource details pages, and in the bottom task and alarm pane. The Alarm Message page displays overall platform alarm information in a dashboard view, making it easier to compare alarms across different dimensions.

Procedure

1. In the navigation pane, choose **O&M Management** > **Alarm Message**.
2. View triggered alarm messages.
   - **Alarm Statistics in Recent 1 Week**: Displays alarm statistics from the last seven days in a bar chart. The sampling interval is 8 hours. Hover over the bar chart to view the number of alarms at different levels.
   - **Alarm Distribution in Recent 1 Week**: Displays the proportion of resource alarms from the last seven days in a bar chart. Hover over the bar chart to view the number of alarms for different resource types.
   - **Alarm Message List**: Displays up to 1,000 alarm messages. You can filter alarm messages by resource type and time.

9.5.5.2 Acknowledge Alarm Messages

Acknowledging an alarm message indicates that the administrator has noticed the alarm.

Procedure

1. In the navigation pane, choose **O&M Management** > **Alarm Message**.
2. In the **Triggered Alarms** list, select an alarm message and click **Acknowledge**.
   - Acknowledged alarm messages are not displayed in **Triggered Alarms**. You can view them in all alarm messages.

- After acknowledgement, if the alarm issue is not handled in time, the alarm continues to trigger and push messages according to its rules. To prevent repeated alarm message pushes, set a silence period as needed.

9.5.5.3 Set a Silence Period

A silence period temporarily pauses the push of a specific alarm message.

Procedure

1. In the navigation pane, choose **O&M Management** > **Alarm Message**.
2. In the **Triggered Alarms** or **All Alarms** list, select an alarm message and click **Actions** > **Set Silence Period**.
3. In the **Set Silence Period** dialog, choose a silence period.

9.5.5.4 Restore Alarm Messages

After you manually restore an alarm message, alarm messages in a silence period can continue to be pushed.

Procedure

1. In the navigation pane, choose **O&M Management** > **Alarm Message**.
2. In the **Triggered Alarms** or **All Alarms** list, select an alarm message and click **Actions** > **Restore Alarm**.

9.6 Tasks, Events, and Audit

NexaSphere records operations, scheduling, HA recovery, and API calls during platform operation through tasks and events. Administrators can use these records to trace operation progress, confirm execution results, and audit user or system access to and changes on platform resources.

9.6.1 View Tasks

The Task page is used to view operation tasks, HA tasks, and scheduling tasks on the platform.

Task Type	Description
Operation Task	Displays completed and ongoing operation task logs. You can view tasks by time range, search by task description or login IP address, filter by task result, sort by start time, end time, or duration, export CSV files, adjust the number of records per page, and cancel,

Task Type	Description
	pause, or resume ongoing tasks. Operation tasks are retained for 90 days by default.
HA Task	Displays all VM HA task logs on the platform. You can view tasks by time range, search by VM name, previous host, or target host, filter by task result, sort by start time or end time, export CSV files, and adjust the number of records per page.
Scheduling Task	Displays VM scheduling logs actively triggered by the management node. You can view tasks by time range, search by migration object or UUID, and adjust the number of records per page. Scheduling tasks are retained for 90 days by default.

9.6.2 View Events

Events continuously monitor and retain platform access and usage activities, including platform operations performed through the console, API services, and developer tools.

Event Type	Description
Resource Action Events	Displays event information for API calls made during platform resource operations. Click an API name to view event details, including basic information, API request, and API response. You can view events by time range, search by API name, resource UUID, resource name, resource type, or operator IP address, filter by task result, export CSV files, and adjust the number of records per page.
Login Operation Events	Displays event information for login and logout API operations. You can view events by time range, search by API name, login IP address, or browser, filter by task result, export CSV files, and adjust the number of records per page.

9.7 Log Collection and Management

Log collection and management enables administrators to collect platform logs by time range and log type and manage or download collected logs for troubleshooting and analysis.

9.7.1 Collect Logs

Prerequisites

- Reserve sufficient storage space. Insufficient capacity will cause log collection to fail.
- The platform can retain up to three collected log packages. To collect new logs, delete collected logs as needed.
- When downloading logs through a URL, make sure the download environment can communicate with the platform management network.

Procedure

1. In the navigation pane, choose **O&M Management** > **Log Collection**.
2. On the **Log Collection** page, click **Collect Log**.
3. Configure the log collection task.

Option	Description
Log Type	Select All Logs or Specified Logs. • All Logs collects operation logs, event logs, management node logs, compute node logs, data storage logs, image storage logs, and database logs. • Specified Logs collects only the selected log types.
Time Range	Specify the time range for collecting logs. A custom time range can be up to 5 days.
Auto Download	Specify whether to automatically download successfully collected logs.

4. Review the configuration and click **OK**.

9.7.2 Manage Collected Logs

Procedure

1. In the navigation pane, choose **O&M Management** > **Log Collection**.
2. Manage collected logs on the **Log Collection** page.
   a) Click **Download** to download a collected log package.
   b) Click **Delete** to delete a single collected log package.
   c) Click **Delete All Logs** to delete all collected log packages.

9.8 SNMP Monitoring and Notification Integration

NexaSphere supports integration with third-party monitoring platforms through SNMP. Third-party platforms can use MIB files to collect resource monitoring data or parse alarm messages. SNMP Trap receivers can also be used as alarm endpoints to receive alarm messages pushed by the platform.

Key SNMP Concepts

SNMP Protocol	A protocol used for managing devices on a network.
Network Management System	A system that monitors and manages network devices through SNMP. In this feature, the network management system corresponds to the third-party monitoring platform.
SNMP Agent	An agent process in managed objects that responds to requests from the network management system or proactively sends Trap information. In this feature, the SNMP Agent role is provided by the NexaSphere management node.
MIB Library	A database maintained by the SNMP Agent that defines managed object attributes. Third-party platforms collect monitoring data or parse alarm messages based on the MIB library.
SNMP Trap Receiver	A third-party server that receives alarm messages from NexaSphere. It can be added as an alarm endpoint and associated with alarms.

Notes

The SNMP components were updated in NexaSphere 1.10.25. If you upgraded from an earlier version and have integrated with an external monitoring platform, download the new MIB file and update it to the external monitoring platform to avoid missing or abnormal monitoring data.

9.8.1 Configure SNMP Management

- Enable SNMP Management.
  a) In the navigation pane, choose **System Management** > **SNMP Management**, click **Enable SNMP Management**.
  b) In the **Enable SNMP Management** dialog, configure the following options.

Option	Description
SNMP Agent Port	Specify the port used to receive and respond to requests from third-party monitoring platforms. The default value is 1160. The value range is 1024 to 65565.
Protocol Version	v2c and v3 are supported. • For v2c, set **Community**. • For v3, set **Username**, and enable **User Authentication** and **Data Encryption** as needed.
Community String	Set a community string for the connection authentication between the third-party monitoring platform and this platform
Username	Set the username used for SNMP v3 access authentication.
User Authentication	Used for secure authentication between the third-party monitoring platform and the platform. After this option is enabled, set the authentication protocol and authentication password.
Data Encryption	Used to encrypt communication messages between the platform and the third-party monitoring platform. After this option is enabled, set the encryption protocol and encryption password.
SNMP Trap Receiver	Enter the third-party server name, IP address, and port for receiving pushed alarm messages.

c) Review the configuration and click **OK**.

- Modify SNMP configuration.

  In the navigation pane, choose **System Management** > **SNMP Management** > **SNMP Configurations**, click **Modify Configuration**, and modify SNMP configuration as needed.

- Disable the SNMP service.

  In the navigation pane, choose **System Management** > **SNMP Management** > **SNMP Configurations**, and click **Disable**.

  - After SNMP is disabled, third-party platforms no longer proactively obtain resource monitoring data from the platform.

- After SNMP is disabled, the platform still pushes alarm messages to SNMP Trap receivers that have been added as endpoints.
- The platform retains the current SNMP configuration for use the next time SNMP is enabled.

9.8.2 Download MIB Files

MIB files are used by third-party platforms to collect NexaSphere resource monitoring data and parse alarm messages.

Procedure

1. In the navigation pane, choose **System Management** > **SNMP Management** > **SNMP Configurations**.
2. On the **SNMP Configurations** tab, click **Download MIB File**.

9.8.3 Manage SNMP Trap Receivers

Procedure

1. In the navigation pane, choose **System Management** > **SNMP Management** > **SNMP Trap Receiver**.
2. Manage SNMP Trap receivers.
   - Click **Add SNMP Trap Receiver** to add a receiver.
   - Select a receiver and click **Actions** > **Modify Configuration** to modify receiver configuration.
   - Select a receiver and click **Actions** > **Delete** to delete the receiver.

   Note:
   If the selected SNMP Trap receiver has been added as an alarm endpoint, the corresponding endpoint is deleted at the same time.

9.9 Resource Tagging and Search

Tags and custom attributes mark resources by service, environment, owner, resource purpose, and other dimensions. Administrators can use tags and custom attributes to quickly filter target resources, improving routine inspection, troubleshooting, and resource governance efficiency.

Tags

Tags are classified into admin tags and user tags. Admin tags are created and owned by administrators. User tags are created and owned by regular users. Administrators can detach

or delete regular users' tags. After a resource owner is changed, user tags on the resource are automatically detached, while administrator tags are not affected.

Resource	Support Tagging	Support Tagging During Creation
Host	Yes	Yes
VM	Yes	Yes
Bare Metal Host	Yes	No

- Changing tag owners is not supported.
- For the same role, only one tag with the same name and color can be created.
- The tag name must be no longer than 20 characters.
- Admin tags can be attached to hosts, VMs, and bare metal hosts.
- User tags can be bound only to VMs.

Custom Attributes

A custom attribute consists of an attribute key and attribute values. After a custom attribute is created, the attribute key is delivered to resources of the selected type. Administrators can set attribute values for specific resources and search resources by attribute key and value.

9.9.1 Manage Tags

- Create a tag.
  a) In the navigation pane, choose **O&M Management** > **Tag**.
  b) On the **Tag Management** page, click **New Tag**.
  c) In the **New Tag** dialog, configure the following options.

Option	Description
Tag Preview	Displays how the tag will appear.
Name	Enter a name for the tag.  **Note:** • Tags must be unique. Users with the same role can create only one tag with the same name and color.

Option	Description
	• Tag names must not exceed 20 characters in length.
Description	Enter a brief desctiption for the tag.
Color	Select color for the tag.

d) Click **OK**.

- Attach or detach tags from the resource perspective.

  In **Inventory** > **Host and VM/Bare Metal Management**, select the target resource, click **Actions** > **Tag and Attribute** > **Tag Management**, and add or remove tags as needed.

- Attach or detach resources from the tag perspective.

  In **O&M Management** > **Tag**, open the tag details page, open the **Associated Resource** tab, select resources, and click **Attach** or **Detach**.

- Delete a tag.

  On the **Tag** page, select a tag, click **Actions** > **Delete**, and confirm the deletion.

9.9.2 Search Resources by Tag

Tags can be used to quickly filter resources bound to tags.

Context

The following steps use VM resources as an example to describe how to filter resources by tag.

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. Select a host from the resource tree.
3. On the host details page, click the **Virtual Machine** tab.
4. On the **Virtual Machine** tab, click **Tag**.
5. Select a tag from the tag drop-down list.
   - Tags are displayed by administrator tag and user tag categories.
   - You can enter a tag name to search for the target tag.
6. Click **OK**.

9.9.3 Create a Custom Attribute

Procedure

1. In the navigation pane, choose **O&M Management** > **Custom Attribute**.
2. On the **Custom Attribute** page, click **New Custom Attribute**.
3. In the **New Custom Attribute** dialog, configure the following options.

Option	Description
Type	Select a resource type. Options include global, virtual machine, host, data storage, distributed switch, distributed port group, and bare metal instance.
Attribute Key	Enter an attribute key name.
Description	Optional. Enter a description for the custom attribute.
Attribute Value	Enter an attribute value name.

4. Click **OK**.

9.9.4 Add Custom Attribute Values to Resources

After you create a custom attribute, the attribute key is automatically associated with all resources of the selected type. You can add attribute values to each resource.

Context

This task uses a VM resource as an example to describe how to set custom attributes for a VM.

Procedure

1. In the navigation pane, choose **Inventory** > **VM and Host**.
2. From the resource tree, select the target VM.
3. On the target VM details page, click **Actions** > **Tag and Attribute** > **Set Custom Attribute**.
4. In the **Set Custom Attribute** dialog, configure attribute values as needed.
   - To add an attribute value to an existing attribute key, select or enter the attribute value in the **Attribute Value** column.
   - To create a new custom attribute, click **Add**, and then enter an attribute key and attribute value.
   - When you configure custom attributes for multiple VMs in batches, the batch configuration overwrites the corresponding attribute values for all selected resources. If a value is left blank, the existing attribute value of each resource is preserved.
5. Click **OK**.

9.9.5 Manage a Custom Attribute

Procedure

1. In the navigation pane, choose **O&M Management** > **Custom Attribute**.
2. On the **Custom Attribute** page, select the target custom attribute.
   - Click **Actions** > **Modify Configuration** to edit the attribute key and description of the custom attribute.
   - Click **Actions** > **Add Attribute Value** to add a new attribute value to the custom attribute.
   - Click **Actions** > **Delete** to delete the custom attribute.
   - Click the name of the custom attribute to open its details page. On the details page, you can manage attribute values in detail, including adding attribute values, deleting attribute values, and viewing associated resources.

9.9.6 Search Resources by Custom Attribute

Custom attributes can be used to quickly filter resources by attribute key and value.

Context

The following steps use VM resources as an example to describe how to filter resources by custom attribute.

Procedure

1. In the navigation pane, choose **Inventory** > **Host and VM**.
2. Select a host from the resource tree.
3. On the host details page, click the **VM** tab.
4. On the **VM** tab, click **Custom Attribute**.
5. Select an attribute key and attribute value from the custom attribute drop-down list.

   If the attribute value is empty, resources that have no value configured for the attribute key are filtered.
6. Click **OK**.

10 Security

10.1 Overview

NexaSphere Security helps administrators manage identity and access, platform access boundaries, certificates, API access credentials, key providers, and trusted VM operation. Administrators can control users and permissions through users, user groups, roles, and Single Sign-On, protect platform access through IP allowlists and blocklists, login policies, password strength, and two-factor authentication, and secure UI access through certificate management.

AccessKeys provide access credentials for third-party platforms that call NexaSphere APIs.
Key providers establish trust between NexaSphere and external KMS systems or the native key service, and provide keys for VM encryption, disk encryption, and vTPM.

10.2 Identity and Access Management

Identity and Access Management provides unified user identity management and access control. Administrators can manage user identities, user groups, roles and permissions, and SSO servers. You can create local users or connect to OIDC, AD, or LDAP authentication servers to synchroniz e users from external authentication systems to the platform.

Roles define the operations that users and user groups can perform. By assigning roles to users or user groups and combining them with resource sharing relationships, administrators can control the scope of platform resources that different users can access.

10.2.1 Prerequisites

Before you use NexaSphere Identity and Access Management, make sure the license meet the following requirements:

- To use Single Sign-On, roles, and user groups, make sure NexaSphere is installed with a valid Advanced Edition license.

10.2.2 User Management

A user represents an individual and is the basic unit in identity and access management. Users are created by admins or synchronized from a unified authentication system, and are managed by admins. By sharing resources with users and assigning roles to them, you can achieve fine-grained control over resource ownership and permissions.

Some key features of users include:

- Users can be either local users or SSO users. Local users are created directly by admins, while SSO users are synchronized from a unified authentication server to the platform.
- User quotas are standards set by admins to control the total amount of resources allocated to users, including compute resources, data storage resources, network resources, and so on.
- Users can be members of one or more user groups.
- Users can be assigned one or more roles. When a user is assigned multiple roles, the user has the combined permissions of those roles. After a user joins a user group, the user inherits the roles assigned to that group in addition to any roles already assigned to the user.

10.2.2.1 Create a User

Create a local user, assign resource ownership and roles, and then use the user to log in.

Prerequisites

- You need to have admin permissions.
- If you need to use the roles and user groups functionality, Make sure NexaSphere is installed with a valid Advanced Edition license.

Procedure

1. In the navigation pane, choose **System Management** > **User Management** > **User**.
2. On the **User** page, click **New User**.
3. In the **New User** dialog, configure the following options:

   Basic Information

Option	Description
Username	Set a username for the regular user, which serves as the unique identifier for logging into the platform. The username cannot be changed after the user is created.
Description	Enter a brief description for the user.
Password	Set a login password for the user. You can adjust the password setting requirements by modifying **Platform Login Password Strength**. For more information, see *Security Setting Options*.
Confirm Password	Re-enter the login password for confirmation.

Option	Description
Role	Assign roles to the user. After roles are assigned, the user has the permissions granted by those roles.
User Group	Add the user to a user group. After joining, the user will inherit all roles and shared resources from the user group.

Share Resource

Option	Description
Share Resource	Specify the resources to be shared with the current user, including virtual machines, images, templates, distributed switches, and distributed port groups.

4. Review the configuration and click **OK**.

10.2.2.2 Disable or Enable a User

After disabling a regular user, the user will not be able to log in to NexaSphere unless the admin enables the user.

Prerequisites

You need to have admin permissions.

Procedure

1. In the navigation pane, choose **System Management** > **User Management** > **User**.
2. On the **User** page, select the target user and then click **Actions** > **Disable**.
3. Review the selected items and click **OK**.
4. To enable the user later, click **Actions** > **Enable**.

10.2.2.3 Modify User Configuration

Edit the user's basic information, such as the roles assigned to the user, the user groups they joined, and the resources shared with the user.

Prerequisites

You need to have admin permissions.

Procedure

1. In the navigation pane, choose **System Management** > **User Management** > **User**.
2. On the **User** page, select the target user and then click **Actions** > **Modify Configuration**.

3. In the **Modify Configuration** dialog, make the necessary changes to the configuration.

10.2.2.4 Change a User Password

Prerequisites

You need to have admin permissions.

Procedure

1. In the navigation pane, choose **System Management** > **User Management** > **User**.
2. On the **User** page, select the target user and then click **Actions** > **Change Password**.
3. In the **Change Password** dialog, enter the new password and confirm it again, then click **OK**.

10.2.2.5 Convert a User to an Admin User

Procedure

1. In the navigation pane, choose **System Management** > **User Management** > **User**.
2. On the **User Management** page, select the target user, and then click **Actions** > **Change to Admin User**.
3. After you read and confirm the risk information, click **OK**.

Result

Note:

- Once a regular user is changed to an admin user, the change cannot be undone.
- After the change, the previously assigned roles, user group memberships, and shared resources will be automatically removed from the regular user.

10.2.2.6 Delete a User

Prerequisites

You need to have admin permissions.

Procedure

1. In the navigation pane, choose **System Management** > **User Management** > **User**.
2. On the **User** page, select the target user and then click **Actions** > **Delete**.
3. In the **Delete User?** dialog, carefully read the risk warnings.

Result

Note:

- The deleted user will no longer be able to log in to the platform, and ownership of their resources will be transferred to admin.
- Deleting an SSO user does not affect the user information in the source authentication server.

10.2.3 User Group Management

A user group is a collection of users that supports permission control at the group level. With user groups, you can assign permissions to multiple users for easier management. For example, if you have a user group named UserGroup-1 and assign storage-related roles to this user group, all users in the group automatically inherit the role permissions from UserGroup-1. If a new user needs storage-related permissions, you can grant the permissions by adding the user to UserGroup-1. When users change, you do not need to modify permissions for each user. You can remove users from the user group as needed.

Some key features of user groups include:

- A user group can contain multiple users, and a user can belong to multiple user groups.
- User groups cannot be nested. User groups can contain only users, not other user groups.
- A user group can be assigned multiple roles. When a user group is assigned multiple roles, users in the group inherit the combined permissions of those roles.

10.2.3.1 Create a User Group

Create a user group, add users to the group, so that all users under the user group can obtain the corresponding permissions, facilitating unified permission management.

Prerequisites

- You need to have admin permissions.
- Make sure NexaSphere is installed with a valid Advanced Edition license.

Procedure

1. In the navigation pane, choose **System Management** > **User Management** > **User Group**.
2. On the **User Group** page, click **New User Group**.
3. In the **New User Group** dialog, set the following parameters:

   Basic Information

Option	Description
Name	Set a name for the user group.
Description	Enter a brief description for the user group.
User	Add users to this user group. After joining, users will inherit all roles and shared resources from this user group.
Role	Assign roles to the user group. After roles are assigned, all users within the group will inherit the permissions granted by these roles.

Share Resource

Option	Description
Share Resource	Share resources with the user group. After sharing, all users within the user group have read access to the shared resources.

4. Review the configuration and click **OK**.

10.2.3.2 Modify User Group Configuration

Edit the basic information of a user group, such as the users within the group, roles assigned to the group, and resources shared with the group.

Prerequisites

- You need to have admin permissions.
- Make sure NexaSphere is installed with a valid Advanced Edition license.

Procedure

1. In the navigation pane, choose **System Management** > **User Management** > **User Group**.
2. On the **User Group** page, select the target user group and then click **Actions** > **Modify Configuration**.
3. In the **Modify Configuration** dialog, make the necessary changes as required.

10.2.3.3 Delete a User Group

Prerequisites

- You need to have admin permissions.
- Make sure NexaSphere is installed with a valid Advanced Edition license.

Procedure

1. In the navigation pane, choose **System Management** > **User Management** > **User Group**.
2. On the **User Group** page, select the target user group and then click **Actions** > **Delete**.
3. In the **Delete User Group?** dialog, carefully read the risk warnings.

Result

Note:
After a user group is deleted, all users within the group will no longer have the roles and shared resources inherited from that group.

10.2.4 Role and Permission Management

A role is a collection of permissions. By assigning roles to users or user groups, administrators grant them the ability to call related APIs and perform resource operations. NexaSphere adopts a Role-Based Access Control (RBAC) authorization model, defining resource permissions based on the user's job function, or role. Through roles, you can achieve fine-grained control over user permissions.

10.2.4.1 Built-in Roles

NexaSphere provides the following built-in roles.

User Type	Role Name	Description
Admin User	System Administrator	Maintains daily system operations.
	Security Administrator	Manages users, security policies, and security attributes.
	Auditor	Manages system event information and auditing.
	Read-Only Role	Has read-only access to system resources without write permissions.
Regular User	VM User	Allows regular users to create VMs and perform basic VM management.

10.2.4.2 Create a Custom Role

To meet diverse access control requirements, you can create custom roles.

Prerequisites

Make sure NexaSphere is installed with a valid Advanced Edition license.

Procedure

1. In the navigation pane, choose **System Management** > **Role**.
2. On the **Role** page, click **New Role**.
3. In the **New Role** dialog, set the following parameters:

   Basic Information

Option	Description
Name	Set a name for the role.
Description	Enter a brief description for the role.

   Permission Configuration

Option	Description
Permission Configuration	Select the UI permissions you want to grant to this role as required. Dependencies may exist between different UI permissions. It is recommended to use platform predefined roles or select all UI permissions.

4. Review the configuration and click **OK**.

10.2.4.3 Clone a Role

To meet diverse access control requirements, in addition to creating custom roles, you can clone existing roles.

Prerequisites

Make sure NexaSphere is installed with a valid Advanced Edition license.

Procedure

1. In the navigation pane, choose **System Management** > **Role**.
2. On the **Role** page, select the target role and then click **Actions** > **Clone Role**.
3. In the **Clone Role** dialog, enter a name and description.
4. Review the configuration and click **OK**.

10.2.4.4 Modify Role Permissions

Edit the interface and API permissions of custom roles.

Prerequisites

- Make sure NexaSphere is installed with a valid Advanced Edition license.
- The selected role is not a system default role.

Procedure

1. In the navigation pane, choose **System Management** > **Role**.
2. On the **Role** page, click the target role name to enter the **Overview** details page.
3. On the **Overview** tab, click the **Edit** icon, and then modify the role's UI permissions as needed.
4. Click **API Permissions** to enter the **API Permissions** tab.
5. On the **API Permissions** tab, click the **Edit** icon, and then modify the role's API permissions as needed.

10.2.4.5 Delete a Role

Prerequisites

- Make sure NexaSphere is installed with a valid Advanced Edition license.
- Make sure the selected role is not a built-in system role.
- Make sure the selected role is not assigned to any user or user group.

Procedure

1. In the navigation pane, choose **System Management** > **Role**.
2. On the **Role** page, select the target role and then click **Actions** > **Delete**.
3. Review the selected items and click **OK**.

10.2.5 Single Sign-On

NexaSphere provides unified identity authentication login services and supports access to unified authentication login systems. SSO users can directly log in to the platform and use platform resources. Currently, OIDC, AD, and LDAP authentication servers are supported.

- **OIDC authentication**: OIDC (OpenID Connect) is an authentication protocol built on the OAuth2 framework. It allows clients to verify user identities and obtain basic user profile information. Through an OIDC authentication server, user information can be synchronized to the platform according to mapping rules, and OIDC users can log in to the platform without a password.
- **AD authentication**: AD (Active Directory) is a directory service for Windows Standard Server, Windows Enterprise Server, and Windows Datacenter Server, providing an independent and standardized login authentication system for increasingly diverse enterprise office applications.

Through an AD authentication server, AD users can be synchronized to the platform and can log in to the platform by using the specified AD login attribute.

- **LDAP authentication**: LDAP (Lightweight Directory Access Protocol) is a protocol for accessing directory services, providing a standardized directory service for increasingly diverse enterprise office applications. Through an LDAP authentication server, LDAP users can be synchronized to the platform and can log in to the platform by using the specified LDAP login attribute.

You can add up to two SSO servers. Only one SSO server of each type can be added.

10.2.5.1 Add an OIDC SSO Server

Prerequisites

- Make sure NexaSphere is installed with a valid Advanced Edition license.
- Make sure fewer than two SSO servers have been added. When the limit is reached, the **Add SSO Server** button is disabled.
- Make sure no OIDC SSO server exists. Only one SSO server of each type can be added.

Procedure

1. In the navigation pane, choose **System Management** > **Single Sign-On**.
2. On the **Single Sign-On** page, click **Add SSO Server**.
3. In the **Add SSO Server** dialog, configure the following options:

   Basic Information

Option	Description
Name	Set a name for the unified authentication server.
Description	Enter a brief description for the unified authentication server.
Type	Select **OIDC**. If an OIDC SSO server already exists, this option is disabled.

   Configuration Information

Option	Description
Client ID	The unique identifier assigned to the platform by the unified authentication system.

Option	Description
Client Secret	The secret key assigned to the platform by the unified authentication system.
Authorization Request URL	The request URL used to obtain authorization under the authorization code grant type.
Token Request URL	The request URL used to obtain an access token from the authentication server.
User Mapping Rules	Establishes the mapping relationship between unified authentication attributes and local attributes, including username. • **Username**: Maps the virtualization platform user name to a specific attribute of users in the unified authentication server.

4. Review the configuration and click **OK**.

10.2.5.2 Add an AD SSO Server

Prerequisites

- Make sure NexaSphere is installed with a valid Advanced Edition license.
- Make sure fewer than two SSO servers have been added. When the limit is reached, the **Add SSO Server** button is disabled.
- Make sure no AD SSO server exists. Only one SSO server of each type can be added.

Procedure

1. In the navigation pane, choose **System Management** > **Single Sign-On**.
2. On the **Single Sign-On** page, click **Add SSO Server**.
3. In the **Add SSO Server** dialog, configure the following options:

   Basic Information

Option	Description
Name	Set a name for the SSO server.
Description	Enter a brief description for the SSO server.
Type	Select **AD**. If an AD SSO server already exists, this option is disabled.

   Server Information

Option	Description
SSL/TLS Encryption	Choose whether to enable SSL/TLS encryption. This option is enabled by default. When enabled, port 636 is used by default and can be changed. When disabled, port 389 is used by default and can be changed.
Primary Server IP/Domain	Enter the primary server IP address or domain and the corresponding port.
Backup Server IP/Domain	Enter the backup server IP address or domain and the corresponding port.

Configuration Information

Option	Description
Base DN	Enter the base DN used to search for AD users' root nodes, defining the scope of synchronized AD users.
User DN	Enter the DN of a special user who has permission to query all users within the base DN scope. This user is used to log in to the AD server and retrieve relevant data.
Password	The password corresponding to the User DN for logging in.
Filter Rule	Enter the filter rule used when synchronizing user information to filter users within the base DN. By default, the `(objectClass=person)` rule is added. **Note:** • Filter rules can be set as single or combined rules, with syntax matching AD filter syntax. • You can control whether the filter acts as a allowlist or blocklist using the `!` symbol. With allowlist filtering, only the user information configured in the filter rules will be synchronized to the platform. With blocklist filtering, user information specified in the filter rules will not be synchronized.

Option	Description
	• The length of filter rules is subject to AD server configuration limits . Exceeding these limits may cause the filter rule to fail, so please confirm in advance.
Login Attribute	Specify the AD user attribute used for logging into the platform.

4. Review the configuration and click **OK**.

10.2.5.3 Add an LDAP SSO Server

Prerequisites

- Make sure NexaSphere is installed with a valid Advanced Edition license.
- Make sure fewer than two SSO servers have been added. When the limit is reached, the **Add SSO Server** button is disabled.
- Make sure no LDAP SSO server exists. Only one SSO server of each type can be added.

Procedure

1. In the navigation pane, choose **System Management** > **Single Sign-On**.
2. On the **Single Sign-On** page, click **Add SSO Server**.
3. In the **Add SSO Server** dialog, configure the following options:

   Basic Information

Option	Description
Name	Set a name for the SSO server.
Description	Enter a brief description for the SSO server.
Type	Select **LDAP**. If an LDAP SSO server already exists, this option is disabled.

   Server Information

Option	Description
SSL/TLS Encryption	Choose whether to enable SSL/TLS encryption. This option is enabled by default. When enabled, port 636 is used by default and can be changed. When disabled, port 389 is used by default and can be changed.

Option	Description
Primary Server IP/Domain	Enter the primary server IP address or domain name and the corresponding port.
Backup Server IP/Domain	Enter the backup server IP address or domain name and the corresponding port.

Configuration Information

Option	Description
Base DN	Enter the base DN used to search for LDAP users' root nodes, defining the scope of synchronized LDAP users.
User DN	Enter the DN of a special user who has permission to query all users within the base DN scope. This user is used to log in to the LDAP server and retrieve relevant data.
Password	The password corresponding to the User DN for logging in.
Filter Rule	Enter the filter rule used when synchronizing user information to filter users within the base DN. By default, the `(objectClass=person)` rule is added. **Note:** • Filter rules can be set as single or combined rules, with syntax matching LDAP filter syntax. • You can control whether the filter acts as a allowlist or blocklist using the ! symbol. With allowlist filtering, only the user information configured in the filter rules will be synchronized to the platform. With blocklist filtering, user information specified in the filter rules will not be synchronized. • The length of filter rules is subject to LDAP server configuration limits . Exceeding these limits may cause the filter rule to fail, so please confirm in advance.
Login Attribute	Specify the LDAP user attribute used for logging into the platform.

4. Review the configuration and click **OK**.

10.2.5.4 Manage SSO Servers

Procedure

1. In the navigation pane, choose **System Management** > **Single Sign-On**.
2. On the **Single Sign-On** page, locate the target SSO server card and click the arrow icon on the right to open the server details.
3. Perform the following operations as needed:
   - If you need to modify the general information, configuration information, or user information mapping rules of the unified authentication server, click **Edit Configuration**.

     Note:
     After modifying the configuration information, unified authentication users who have been synchronized to the platform may no longer be able to log in without a password.
   - To edit the name and description of the unified authentication server, select **More Actions** > **Edit Name and Description**.
   - To delete the unified authentication server, select **More Actions** > **Delete**.

     Note:
     Deleting the unified authentication server will also remove related existing unified authentication user information from the platform, while users in the source unified authentication server remain unaffected.

10.3 Platform Access Security

Platform access security controls user login and platform access. Administrators can restrict access sources through IP allowlists and blocklists, and manage login policies, password strength, two-factor authentication, VNC console passwords, VM password strength, and host password encrypted storage through security settings.

10.3.1 IP Allowlist and Blocklist

IP allowlists and blocklists control platform access by identifying and filtering visitor IP addresses. A blocklist denies access from specified IP addresses, while an allowlist permits access from specified IP addresses.

- If no IP blocklists or allowlists are added, requests from all IP addresses are allowed by default.

- If only an IP blocklist is added, IP addresses in the blocklist are denied access to the platform, while other IP addresses are allowed.
- If the same IP address is added to both lists, the allowlist takes precedence over the blocklist, and requests from that IP address are allowed.
- You cannot use an IP allowlist alone. Add at least one IP blocklist before using an IP allowlist. Otherwise, the IP allowlist does not take effect.

10.3.1.1 Add an IP Allowlist or Blocklist Rule

Procedure

1. In the navigation pane, choose **System Management** > **IP Allowlist/Blocklist**.
2. Click **Add IP Allowlist/Blocklist**.
3. In the **Add IP Allowlist/Blocklist** dialog, configure the following options:

Option	Description
Name	Enter a name for the IP allowlist or blocklist.
Description	Enter a brief description for the IP allowlist or blocklist.
Type	Select blocklist or allowlist.
IP Address	Enter IP addresses, IP address ranges, or IP/mask values. Separate multiple IP addresses with commas. You can add up to 100 entries.

10.3.1.2 Manage IP Allowlist and Blocklist Rules

Procedure

1. In the navigation pane, choose **System Management** > **IP Allowlist/Blocklist**.
2. Select a target list and click Actions.
   - To modify the name and description of the list, select **Edit Name and Description**.
   - To modify the IP addresses in the list, select **Modify Configuration**.
   - To remove the IP access restriction from the platform, select **Delete**.

10.3.2 Security Settings

NexaSphere provides security settings for scenarios that require stronger platform access protection. Administrators can configure platform login policies, password strength, two-factor authentication, VNC console password requirements, VM password strength, and host password encrypted storage.

10.3.2.1 Modify Security Settings

Procedure

1. In the navigation pane, choose **System Management** > **Security Settings**.
2. Select the item to modify, and then click the Edit icon to complete the configuration.

10.3.2.2 Security Setting Options

Category	Item Name	Item Description
Login Policy	Restrict Concurrent Connections for Users	Specify whether to limit the number of simultaneous connection sessions for each user. Default: false. If enabled, each user can only have one active connection session with the platform. When a new session is established, the previous session will be automatically closed.
	Session Timeout	If a session times out, you will need to log in again to access the system. Default: 2 hours. Units: seconds, minutes, hours, and days.
	Login Verification Code Settings	Specify whether to require a verification code after a specified number of consecutive failed login attempts . Default: false. If enabled, you can set the maximum number of continuous login failures that will trigger the verification code requirement. Enter an integer that is equal to or greater than 3.
	Login Password Update Cycle	Specify whether to enforce regular updates of login passwords. Default: false. If enabled, you can specify the update interval. The default interval is 90 days. When the update period is reached, you will be prompted to change your password upon logging into the platform. You can also set the number of recent passwords that cannot be repeatedly used. For example, if you set this parameter to 3, you will not be able to reuse any of your last three passwords when logging into the platform.
	Failed Login Lockout Settings	Specify whether to lock a user's account after consecutive failed login attempts. Default: false. If enabled, set **Max Login Attempts** to an integer from 1 to 10 and **Lockout Duration** to an integer from 0 to 1,440 minutes. When the number of consecutive failed login attempts reaches the limit, the account is locked for the specified duration.

Category	Item Name	Item Description
	Login Password Strength	Specify whether to enforce the user-created login passwords to meet a certain level of complexity and security. Default: false. If enabled, you can specify the required range for password length and enforce the use of a combination of digits, lowercase letters, uppercase letters, and special characters.
	Two-Factor Authentication	Specify whether to enable two-factor authentication (2FA) for platform logins. Default: false.
Virtual Machine	VNC Console Password	Specify whether to require a password for VNC console login. Default: false. If enabled, you can specify the length range for the VNC console password. The default range is 6-8 characters. Additionally, you can enforce the use of a combination of digits, letters, and special characters.
	VM Password Strength	Specify whether to require login passwords for virtual machines. Default: false. 1. The VM password format is m-n, with values ranging from 8 to 32 integers. Default: 8 to 18. The password supports a combination of digits, letters, and special characters. 2. Before you set a login password for a virtual machine, make sure that cloud-init is installed in the VM system image. We recommend that the version of cloud-init be 0.7.9, 17.1, 19.4, 19.4, or later.
Host	Host Password Encryption Policy	Specifies whether and how to encrypt the login password of hosts in the database. Default: None. Options include None and LocalEncryption. • None: Do not encrypt the login password of hosts. • LocalEncryption: Encrypt the login password of hosts by using the encryption feature provided by the platform.

10.4 SSL Certificate

SSL Certificate is used to maintain the HTTPS access certificate for the NexaSphere UI. Administrators can import a third-party certificate, import a system self-signed certificate, update an existing certificate, and switch to HTTP login for the UI in specific scenarios.

10.4.1 Import a Third-Party Certificate

Prerequisites

- You have deployed the latest NexaSphere environment. For a dual-management node environment, ensure that each management node is working properly.
- You need admin permissions to configure the certificates.
- You hold a valid commercial CA-issued certificate.
- Certificate files and certificate chains are supported in CTR or PEM format only. Private keys for certificates must be in KEY or PEM format.

> **Note:**
> If your certificate does not meet these format requirements, convert it accordingly.

Procedure

1. In the navigation pane, choose **System Management** > **SSL Certificate**.
2. On the **SSL Certificate** page, click **Import Certificate**.
3. In the **Import Certificate** dialog, configure the following options:

Option	Description
Import Mode	Select **Third-party Certificate**.
Certificate File	Import or enter the certificate public key. **Note:** • Only CTR and PEM formats are supported. • The certificate content must begin with `----BEGIN CERTIFICATE----` and end with `----END CERTIFICATE----`.
Certificate Private Key	Import or enter the certificate private key. **Note:** • Only KEY and PEM formats are supported. • The private key content must begin with `----BEGIN (RSA/EC) PRIVATE KEY----` and end with `----END (RSA/EC) PRIVATE KEY----`.

Option	Description
Certificate Chain	Import or enter the certificate chain.  **Note:** • Only CTR and PEM formats are supported. • The certificate chain content must begin with `----BEGIN CERTIFICATE----` and end with `----END CERTIFICATE----`.
HTTP Redirection	Optional, enabled by default. When enabled, the system automatically redirects requests from port 80 of the HTTP address to port 443 of the HTTPS address.

4. Review the certificate information and click **OK**.

Result

After successfully importing the third-party certificate, the system will re-establish the session and reconnect to the UI management interface through port 443 of the HTTPS protocol.

10.4.2 Import a System Self-Signed Certificate

Prerequisites

- The latest NexaSphere environment has been deployed. In a dual-management-node environment, make sure each management node works properly.
- You need admin permissions to configure certificates.

Procedure

1. In the navigation pane, choose **System Management** > **SSL Certificate**.
2. On the **SSL Certificate** page, click **Import Certificate**.
3. In the **Import Certificate** dialog, configure the following options:

Option	Description
Import Mode	Select **System Self-Signed Certificate**.
Validity Period	Options include 3 months, 1 year, 3 years, 5 years, or 10 years. The default is 3 years.

Option	Description
HTTP Redirection	Optional, enabled by default. When enabled, the system automatically redirects requests from port 80 of the HTTP address to port 443 of the HTTPS address.
Custom Information	Optional, disabled by default. When enabled, you can customize the system-signed certificate information.

Custom Information

Option	Description
Common Name (CN)	Optional. Set the common name. The default is localhost. The length should be 1 to 64 characters. Only uppercase and lowercase English letters, numbers, and the following special characters are supported: ~`@#$%^&*()-_+={}[]|:;'<>.?/
Organization (O)	Optional. Set the organization name. The default is localhost. The length should be 1 to 64 characters. Only uppercase and lowercase English letters, numbers, and the following special characters are supported: ~`@#$%^&*()-_+={}[]|:;'<>.?/
Organizational Unit (OU)	Optional. Set the department. The length should be 1 to 64 characters. Only uppercase and lowercase English letters, numbers, and the following special characters are supported: ~`@#$%^&*()-_+={}[]|:;'<>.?/
Country/Region (C)	Optional. Set the country/region. Only CN is supported.
State/Province (S)	Optional. Set the state/province. The length should be 1 to 128 characters. Chinese characters, uppercase and lowercase English letters, numbers, and the following special characters are supported: ~`@#$%^&*()-_+={}[]|:;'<>.?/
Locality (L)	Optional. Set the city. The length should be 1 to 128 characters. Chinese characters, uppercase and lowercase English letters, numbers, and the following special characters are supported: ~`@#$%^&*()-_+={}[]|:;'<>.?/
Email Address	Optional. Set the email address.

4. Review the certificate information and click **OK**.

Result

After successfully importing the system-signed certificate, the system will re-establish the session and reconnect to the UI management interface through port 443 of the HTTPS protocol.

10.4.3 Update a Certificate

Prerequisites

- The latest NexaSphere environment has been deployed. In a dual-management-node environment, make sure each management node works properly.
- You need admin permissions to configure certificates.
- If an added certificate has changed or is nearing its expiration date, you need to update the certificate information promptly.

Procedure

1. In the navigation pane, choose **System Management** > **SSL Certificate**.
2. In the **SSL Certificate** page, click **Import New Certificate**.
3. In the **Import Certificate** dialog, update the certificate configuration information.

   Note:
   When updating a certificate, the system checks the current certificate path and writes the certificate information to that path.
4. Review the certificate information and click **OK**.

Result

After successfully updating the certificate, you can continue to access the UI management interface through port 443 of the HTTPS protocol.

10.4.4 Switch to HTTP Login for the UI

Prerequisites

- The latest NexaSphere environment has been deployed. In a dual-management-node environment, make sure each management node works properly.
- You need admin permissions to configure certificates.
- You have configured the SSL certificate.

Procedure

1. In the navigation pane, choose **System Management** > **SSL Certificate**.

2. On the **SSL Certificate** page, click **Switch to HTTP**.
3. In the confirmation dialog, review and confirm the risk warning information.

Result

After successfully switching to HTTP protocol for accessing the UI management interface, the system will re-establish the session and reconnect to the UI management interface through port 80 of the HTTP protocol.

10.5 AccessKey Management

An AccessKey is a security credential that authorizes third-party platforms to call NexaSphere APIs and access virtualization resources. It includes an AccessKey ID and an AccessKey Secret. Keep AccessKey credentials secure, and enable, disable, or delete them as needed.

Permission Relationships

- Administrators can create multiple AccessKeys, while regular users can create up to two AccessKeys.
- Administrators can delete AccessKeys created by regular users.
- An AccessKey has full permissions of its creator.

10.5.1 Generate an AccessKey

Procedure

1. In the navigation pane, choose **System Management** > **AccessKey Management**.
2. Click **Generate AccessKey**.

10.5.2 Manage AccessKeys

Procedure

1. In the navigation pane, choose **System Management** > **AccessKey Management**.
2. Select the target AccessKey.
   - To enable or disable the AccessKey, click **Enable** or **Disable**.
   - If you no longer need to access virtualization resources by using APIs, click **Delete** to delete the AccessKey.

10.6 Key Provider Management

Key providers supply keys for VM encryption, disk encryption, and vTPM. NexaSphere supports Standard Key Providers and Native Key Providers. A Standard Key Provider requires two-way

trust with an external KMS. A Native Key Provider is provided by the platform and can be used without a third-party KMS.

When you create an encrypted resource or change the encryption state of a resource, the platform obtains keys from the default key provider. Before you use encryption, make sure the default key provider is connected and its key service is available. Back up the Native Key Provider before use so that you can restore the key service after a platform failure or rebuild.

10.6.1 Configure and Manage a Standard Key Provider

10.6.1.1 Add a Standard Key Provider

You can add an external key server (KMS) to support various encryption capabilities of the system, thereby enhancing overall system security. The KMS must support the Key Management Interoperability Protocol (KMIP) 1.1 or later.

Prerequisites

Make sure the platform is installed with a valid Advanced Edition license.

Procedure

1. In the navigation pane, select **Inventory** > **VM and Host**.
2. Select the root node.
3. On the root node details page, click **Key Providers**.
4. On the **Key Providers** tab, click **Add Key Provider**.
5. In the **Add Key Provider** dialog, configure the following options:

Option	Description
Name	Enter a name for the key provider.  **Note:** The name of each key provider serves as its identifier. Ensure the name is globally unique.
Description	Enter a brief description for the key provider.
Type	Select **Standard Key Provider**.
IP Address/ Domain Name	Enter the IP address or domain name of the KMS.

Option	Description
Port	Enter the connection port of the KMS. The default port is 5696.
Password Protection	Some KMS systems allow you to isolate encryption keys used by different users or groups by specifying a username and password. Enable this option and enter the corresponding username and password.

6. After you confirm the configuration details are correct, click **OK**.
7. Click **Trust**.

   After the key provider is added successfully, its status shows as "Connected | Not Trusted by KMS".

What's next

- After you add a standard key provider, you must establish two-way trust for it to work properly. For more information, see *Use KMS Certificate and Private Key to Make KMS Trust NexaSphere*, *Use the New Certificate Signing Request (CSR) to Make KMS Trust NexaSphere*, and *Make NexaSphere Trust the KMS*.
- Scenario practices: *Use the Thales Key Provider*, *Use the HashiCorp Key Provider*

10.6.1.2 Establish Trust

10.6.1.2.1 Use KMS Certificate and Private Key to Make KMS Trust NexaSphere

Some KMS vendors require you to upload the KMS server certificate and private key to NexaSphere.

Procedure

1. Navigate to the **Key Providers** page.
2. Locate the standard key provider with which you want to establish trust, and click **Actions** > **Establish Trust** > **Make KMS Trust the Platform**.
3. In the **Make KMS Trust the Platform** dialog, select **KMS Certificate and Private Key** as the trust method.
4. Click **Next**.
5. Upload the certificate and private key files you received from the KMS, or paste the certificate and private key content into the text box.
6. Click **OK**.

What's next

To establish two-way trust, see *Make NexaSphere Trust the KMS*.

10.6.1.2.2 Use the New Certificate Signing Request (CSR) to Make KMS Trust NexaSphere

Some KMS vendors require NexaSphere to generate a Certificate Signing Request (CSR) and submit the CSR to the KMS. The KMS signs the CSR and returns a signed certificate. You must upload the signed certificate back to NexaSphere to establish trust.

Procedure

1. Navigate to the **Key Providers** page.
2. Locate the standard key provider with which you want to establish trust, and click **Actions** > **Establish Trust** > **Make KMS Trust the Platform**.
3. In the **Make KMS Trust the Platform** dialog, select **New Certificate Signing Request (CSR)** as the trust method.
4. Click **Next**.
5. Click **Copy** or **Download** to copy the entire content from the text box or save it as a file.

   Use **Generate New CSR** only when you explicitly need to create a new CSR.
6. Click **OK**.
7. Log in to the KMS and submit the CSR according to the KMS requirements.
8. After you receive the certificate signed by the KMS, click **Actions** > **Establish Trust** > **Upload Signed CSR Certificate** again.
9. Paste the signed certificate into the text box, or click **Upload File** to upload it.

What's next

To establish two-way trust, see *Make NexaSphere Trust the KMS*.

10.6.1.2.3 Make NexaSphere Trust the KMS

If you skipped the step of trusting the KMS when you added the standard key provider, you must establish two-way trust after the KMS trusts NexaSphere.

Procedure

1. Navigate to the **Key Providers** page.
2. Locate the standard key provider with which you want to establish trust, and click **Actions** > **Establish Trust**.
3. Select one of the following options from the Establish Trust menu.

- **Make the Platform Trust KMS**: In the **Make the Platform Trust KMS** dialog, click **Trust**.
- **Upload KMS Certificate**: In the **Upload KMS Certificate** dialog, upload a file or paste the certificate content into the text box. Then click **OK**.

10.6.1.3 Delete a Standard Key Provider

Prerequisites

- When multiple key providers exist, the default key provider cannot be deleted. Set another key provider as the default first, and then perform the deletion.
- To preserve encrypted resources that depend on this key provider, perform a global **Rekey** operation before deletion. The operation uses a new KEK from the current default key provider to rewrap the DEKs of the encrypted resources.

Procedure

1. In the navigation pane, select **Inventory** > **VM and Host**.
2. Select the root node.
3. On the root node details page, click **Key Providers**.
4. On the **Key Providers** tab, select the standard key provider you want to delete, and then click **Actions** > **Delete**.

   Note:

   After you delete a key provider, resources that depend on it, such as vTPM-enabled VMs, encrypted VMs, and encrypted disks, are locked and become unavailable.
5. Read the risk warning, enter the confirmation information, and click **OK**.

10.6.1.4 Use the Thales Key Provider

Context

1. *Add the Thales key provider in NexaSphere*
2. *Create a localhost user in Thales and add it to the admin group*
3. *(Optional) Use the **New Certificate Signing Request (CSR)** method to make Thales trust NexaSphere*
4. *(Optional) Use the **KMS certificate and private key** method to make Thales trust NexaSphere*

Procedure

1. Add the Thales key provider in NexaSphere.

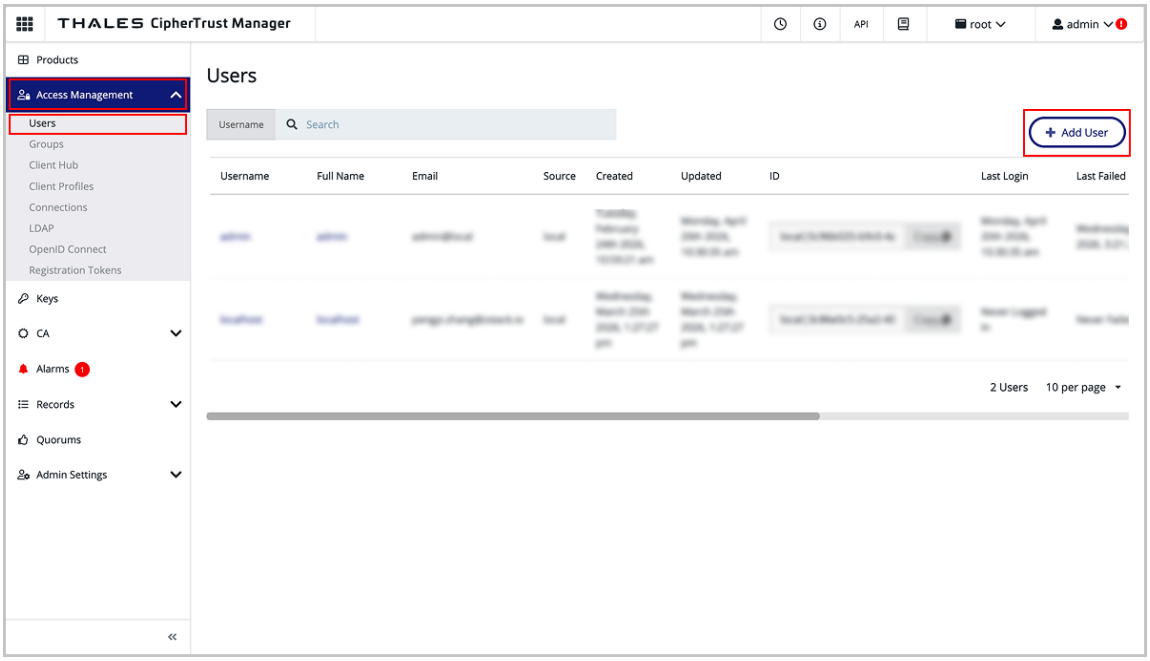
For detailed steps, see *Add a Standard Key Provider*.

2. Create a localhost user in Thales and add it to the admin group.

   a) Log in to Thales.

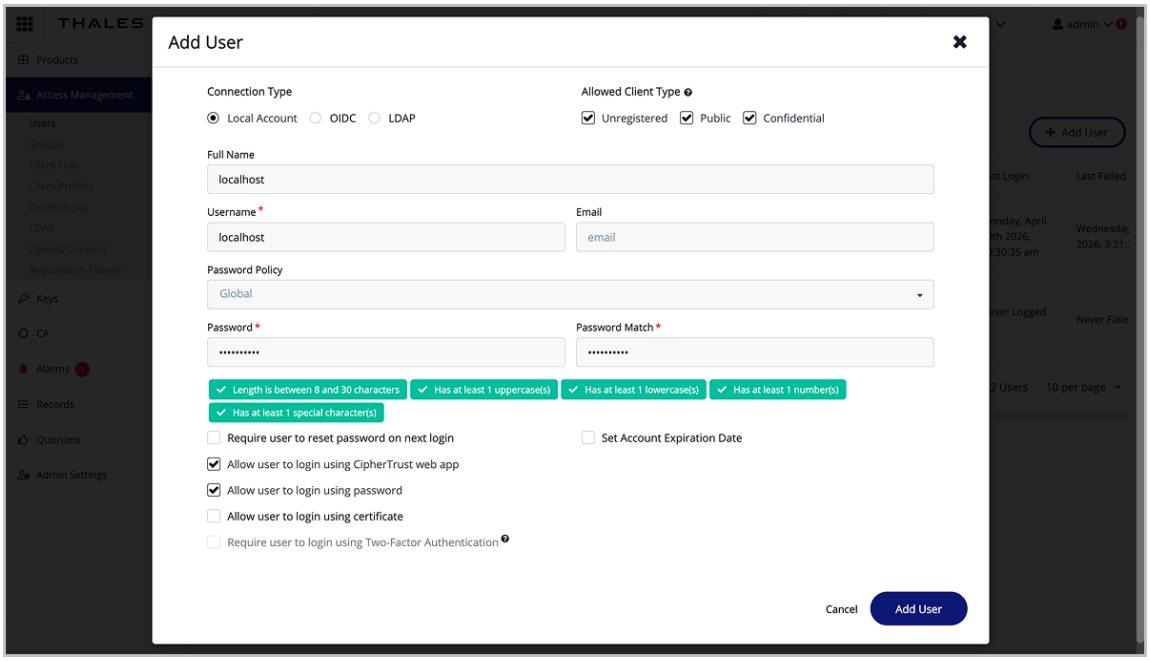
   b) Click **Access Management** > **Users** > **Add User**.

Figure 10-1: Add a User



   c) Create the localhost user.

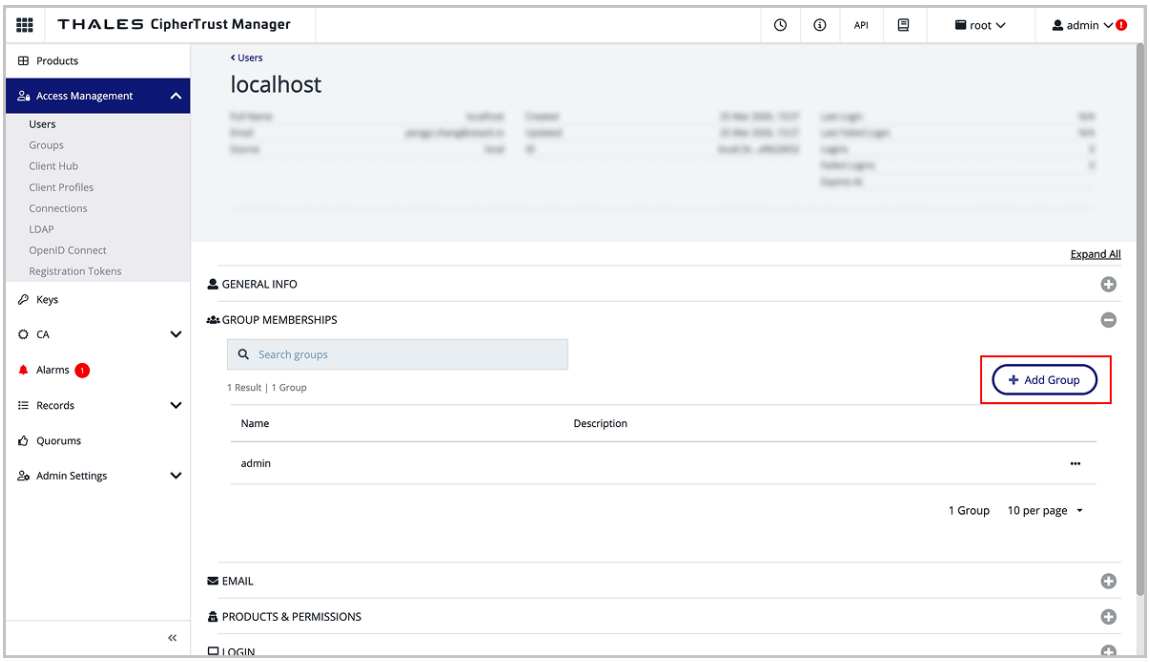
Figure 10-2: Create the localhost User



d) Click the user name to go to its details page.

e) Expand the **GROUP MEMBERSHIP** section, and then click **Add Group**.

Figure 10-3: Add the User to a Group



f) In the **Add member to group(s)** dialog, select **admin**, and then click **Add Group**.

3. (Optional) Use the **New Certificate Signing Request (CSR)** method to make Thales trust NexaSphere.

   a) Log in to NexaSphere.

   b) Navigate to the **Key Providers** page.

   c) Locate the Thales standard key provider you added, and click **Actions** > **Establish Trust** > **Make the KMS Trust NexaSphere**.

   d) In the **Make the KMS Trust NexaSphere** dialog, select **New Certificate Signing Request (CSR)** as the trust method.

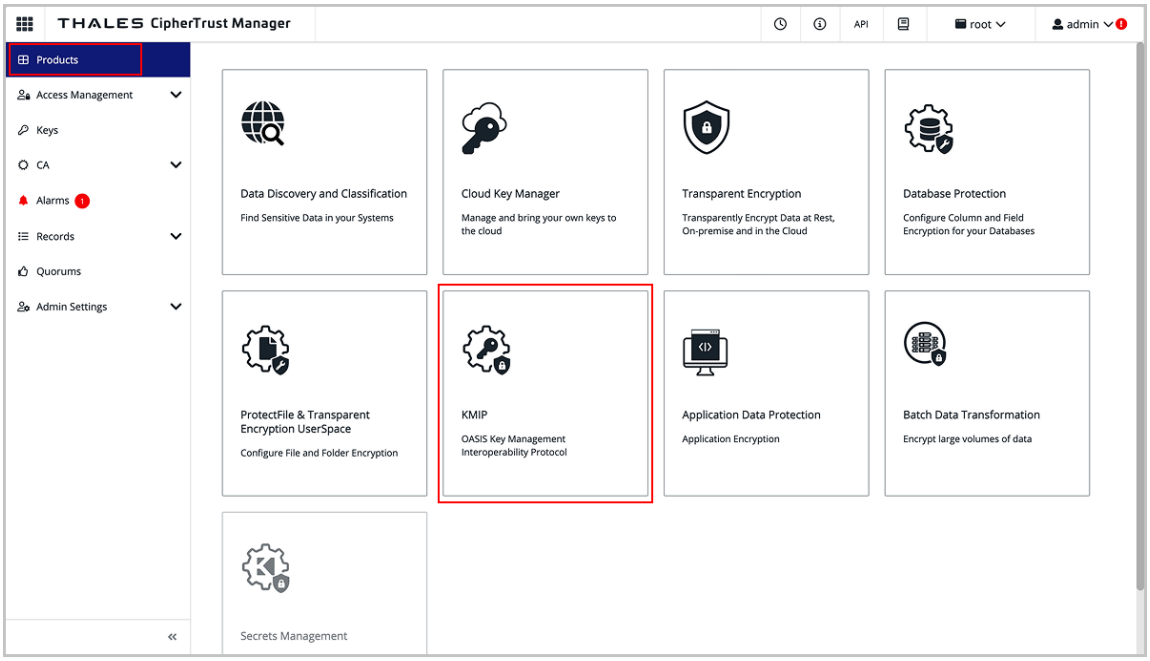
   e) Click **Next**.

   f) Copy the entire CSR content.

   g) Log in to Thales.

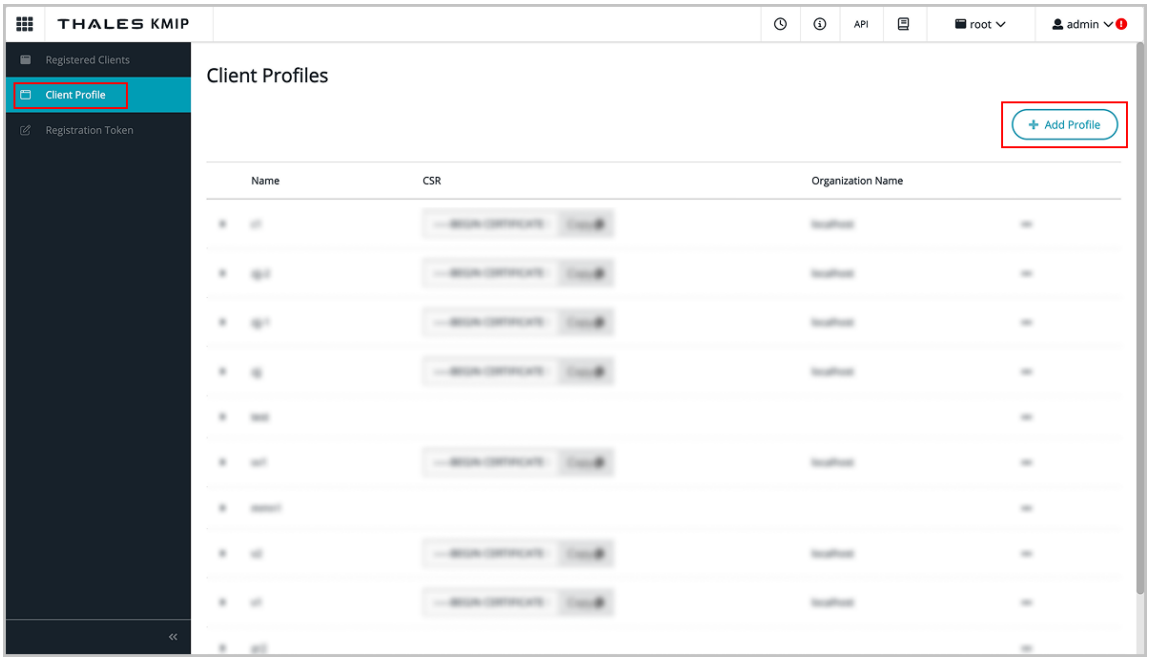
   h) Click **Products** > **KMIP**.

Figure 10-4: Access the KMIP Module



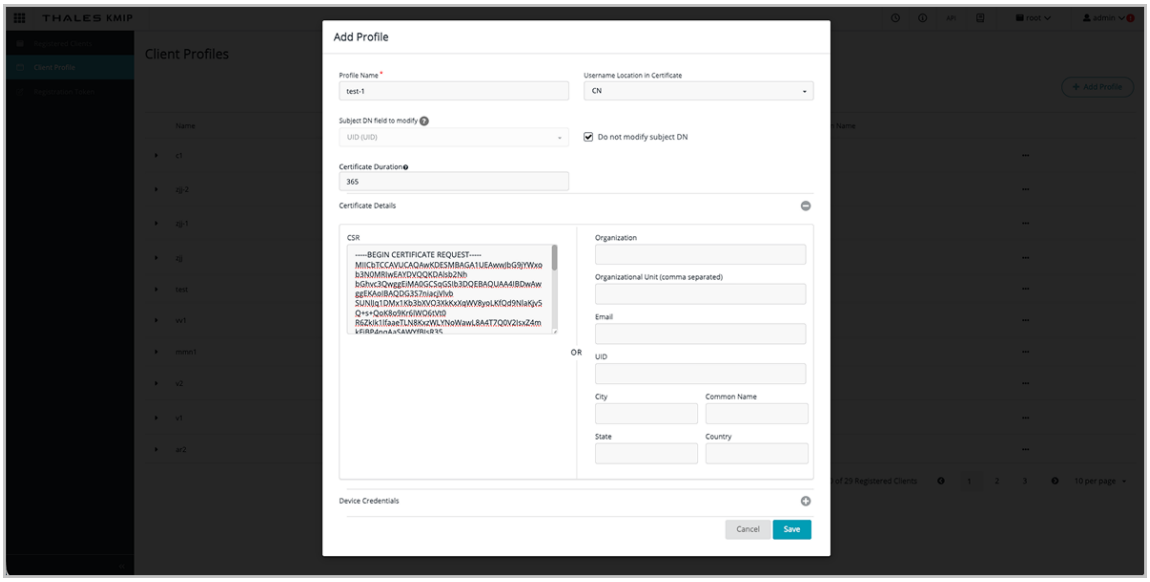
   i) Click **Client Profile** > **Add Profile**.

Figure 10-5: Add a Profile



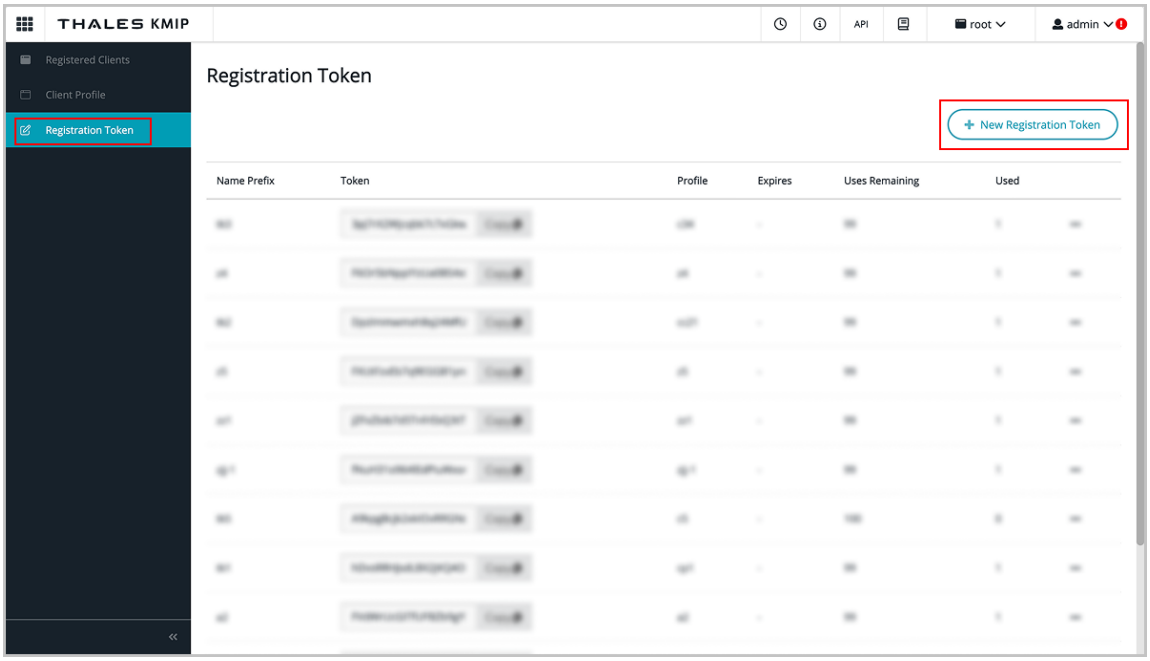
j) In the **Add Profile** dialog, set the profile basic information, paste the CSR information, and then click **Save**.

Figure 10-6: Configure the Profile



k) Click **Registration Token** > **New Registration Token**.

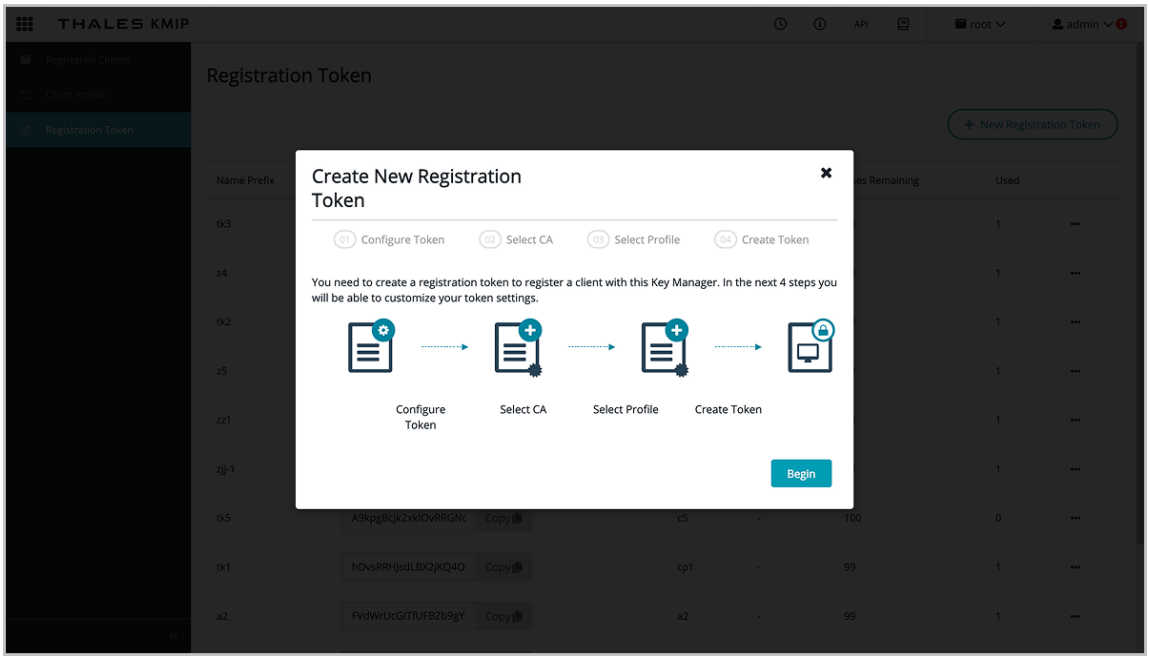
Figure 10-7: New Registration Token



l) In the **Create New Registration Token** dialog, follow the wizard to complete the configuration.

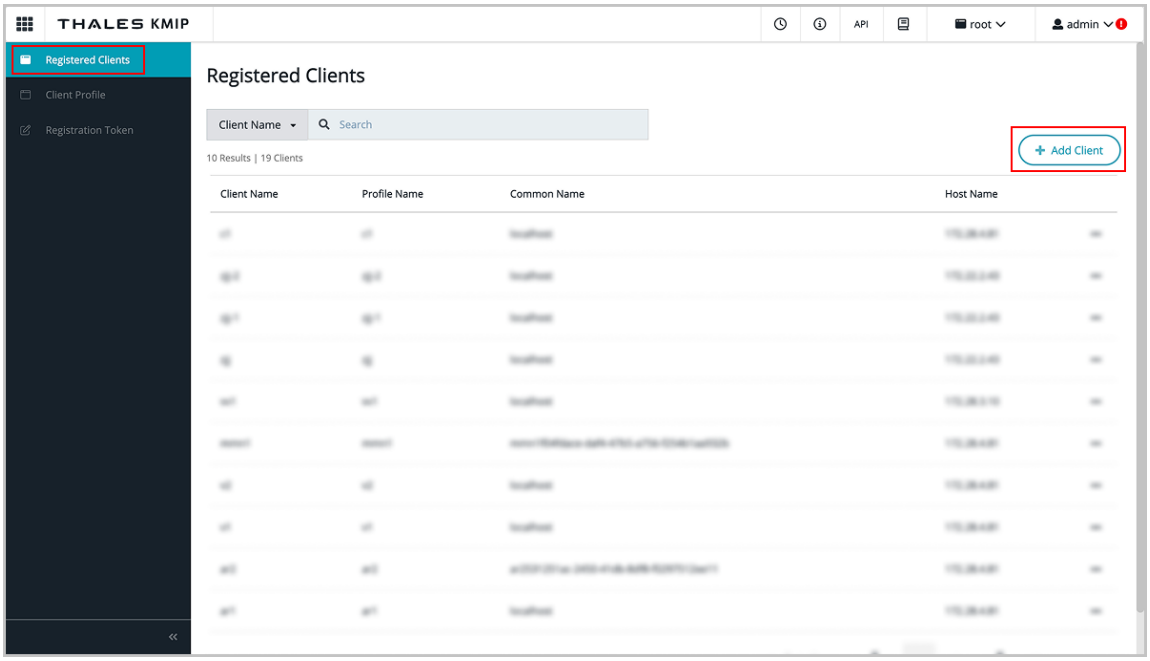
- In step 3 **Select Profile**, select the profile you created earlier.
- In step 4 **Create Token**, copy the generated token.

Figure 10-8: Create a Token



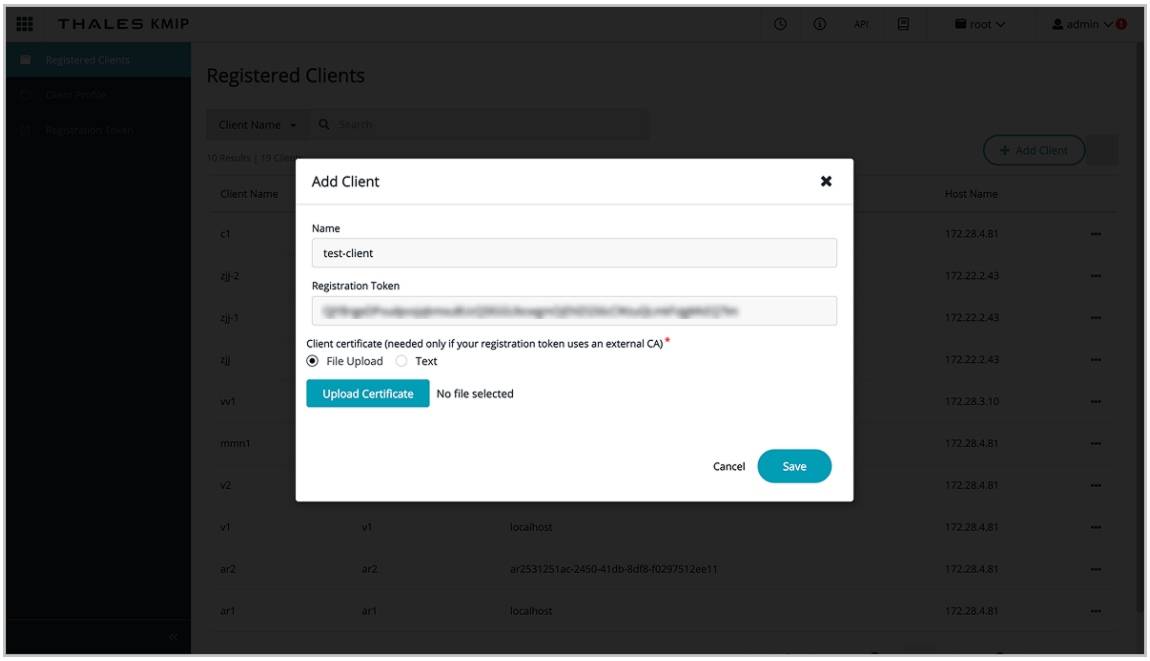
m) Click **Registered Clients** > **Add Client**.

Figure 10-9: Add a Client



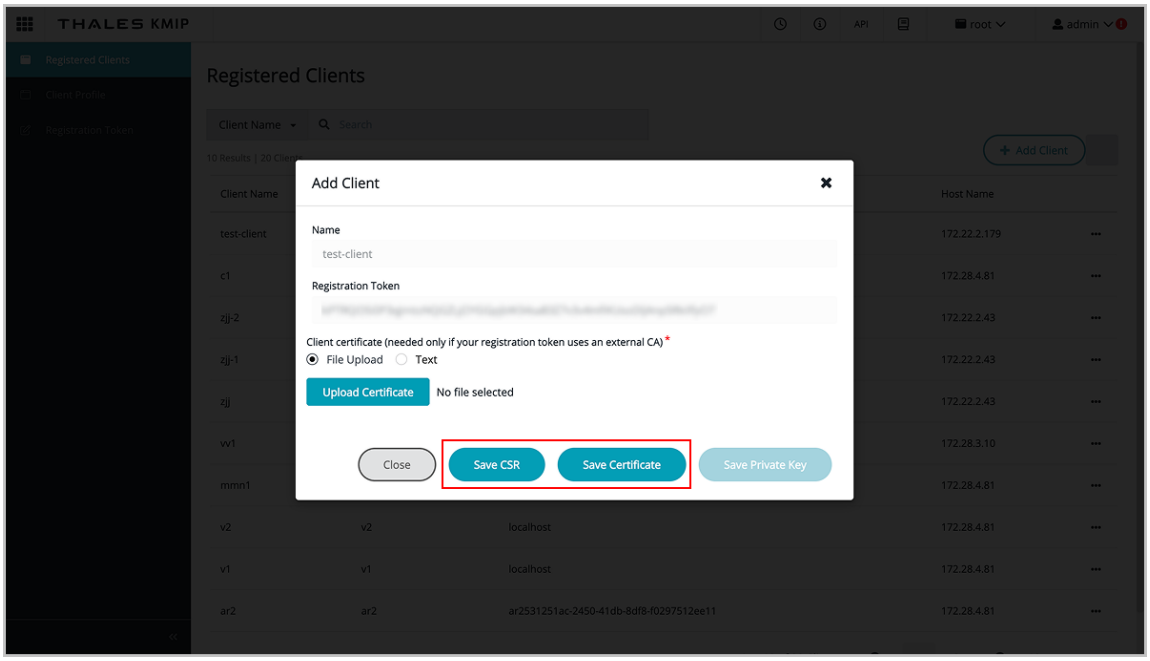
n) In the **Add Client** dialog, set the client name, paste the registration token, and then click **Save**.

Figure 10-10: Configure the Client



o) After saving successfully, click **Save CSR** and **Save Certificate** to download the CSR and certificate.

Figure 10-11: Download the CSR and Certificate



p) Log in to NexaSphere.

q) Navigate to the **Key Providers** page, and then click **Actions** > **Establish Trust** > **Upload Signed CSR Certificate**.

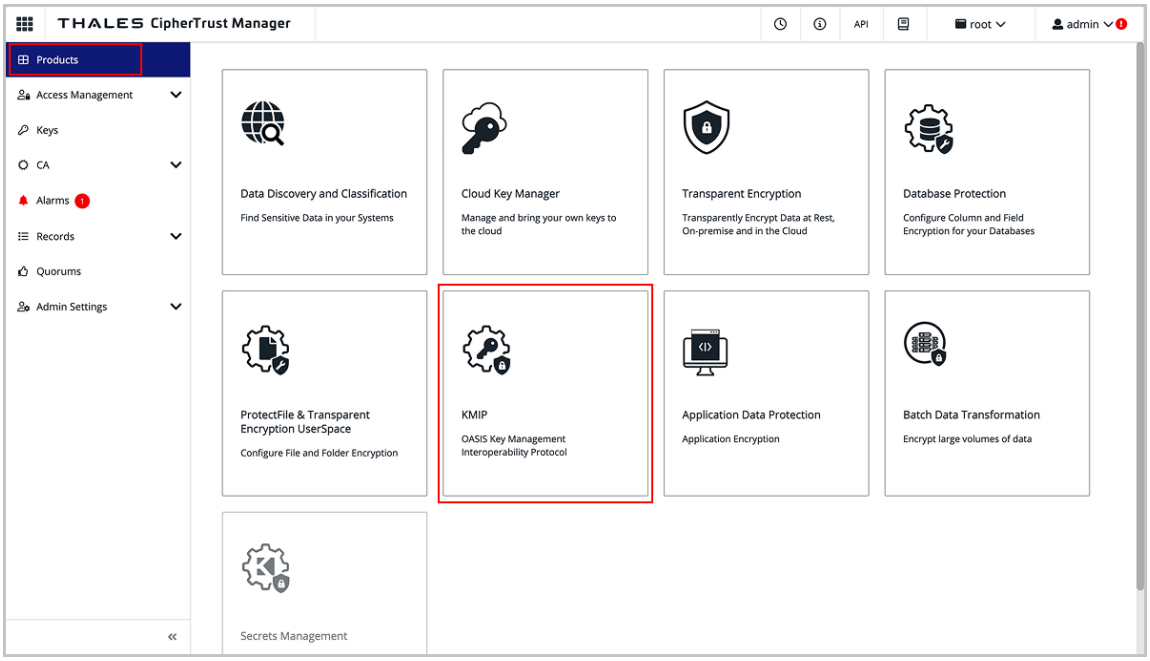
r) Upload the certificate you downloaded from Thales.

4. (Optional) Use the **KMS certificate and private key** method to make Thales trust NexaSphere.

a) Log in to Thales.

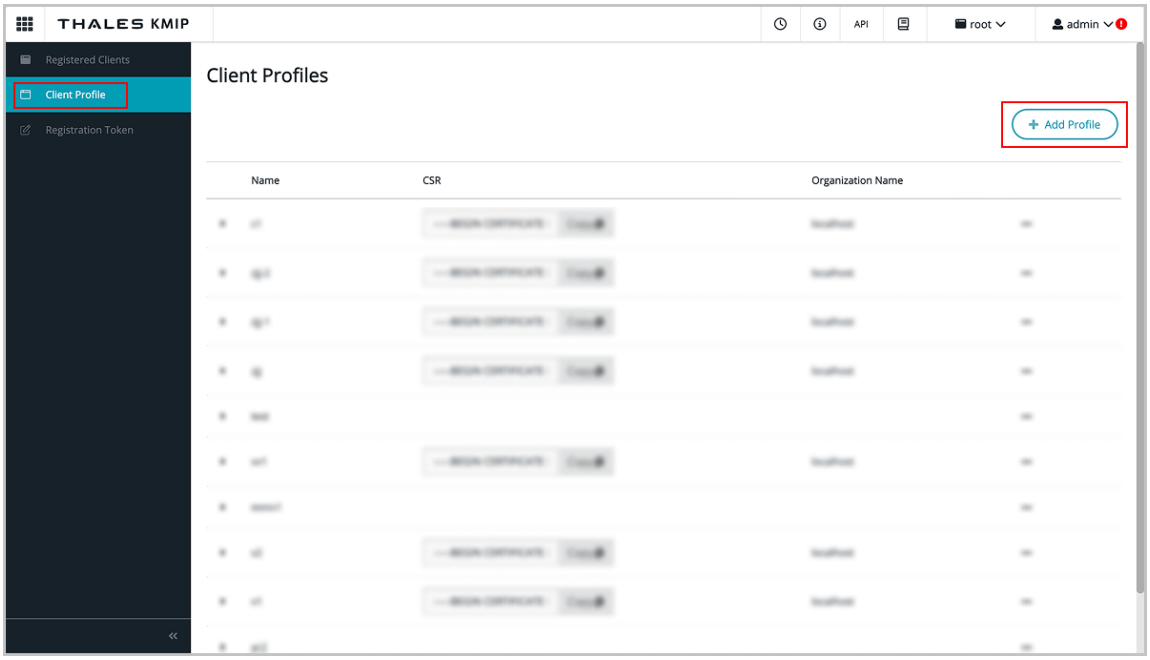
b) Click **Products** > **KMIP**.

Figure 10-12: Access the KMIP Module



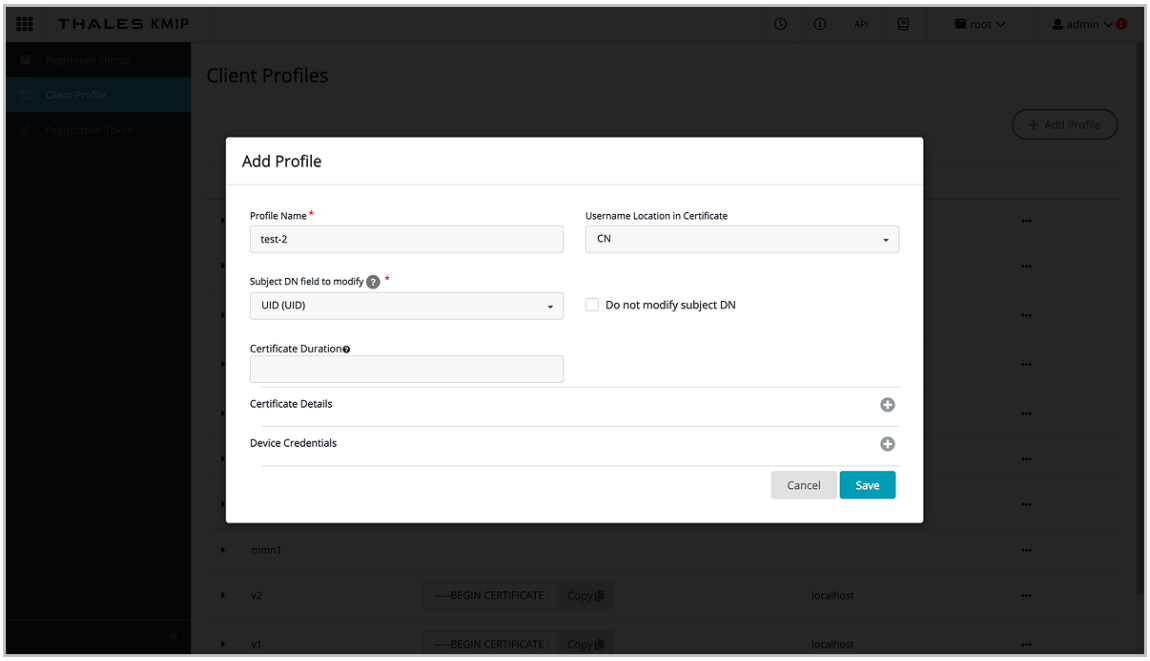
c) Click **Client Profile** > **Add Profile**.

Figure 10-13: Add a Profile



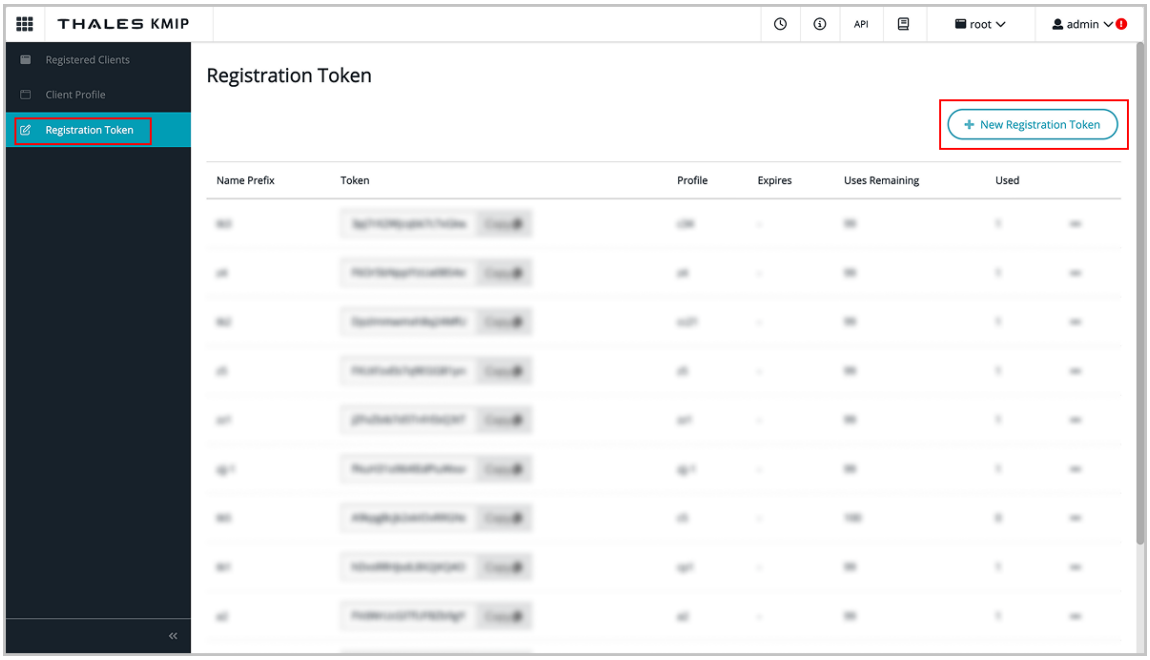
d) In the **Add Profile** dialog, set the profile name, and then click **Save**.

Figure 10-14: Add a Profile



e) Click **Registration Token** > **New Registration Token**.

Figure 10-15: New Registration Token

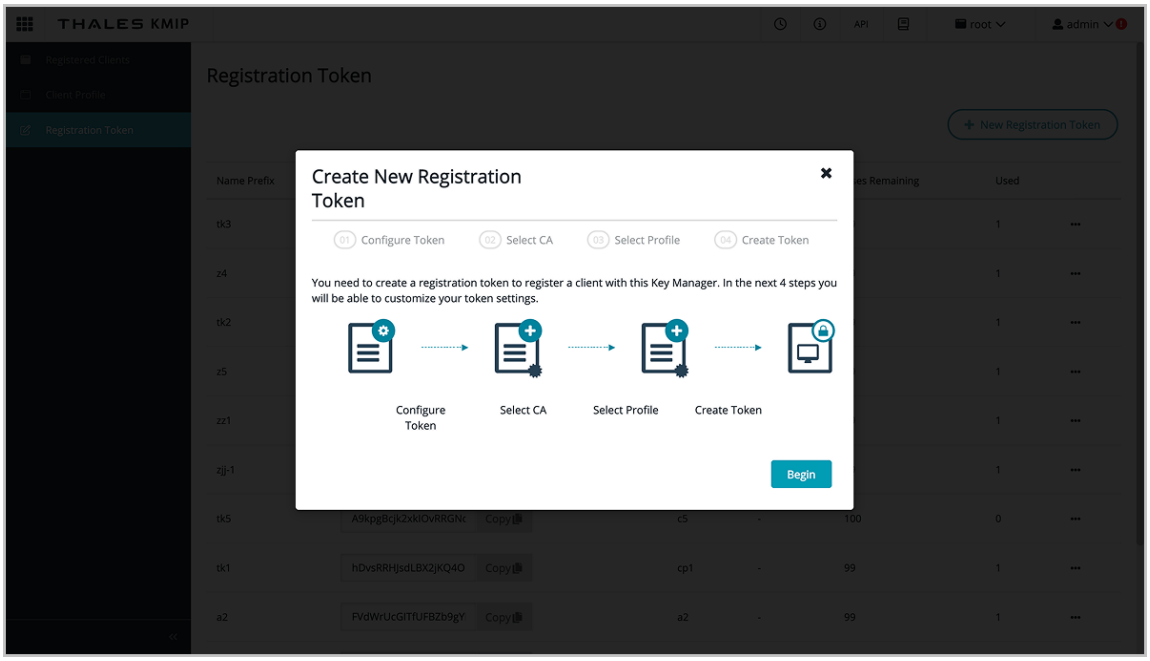


f) In the **Create New Registration Token** dialog, follow the wizard to complete the configuration.

- In step 3 **Select Profile**, select the profile you created earlier.

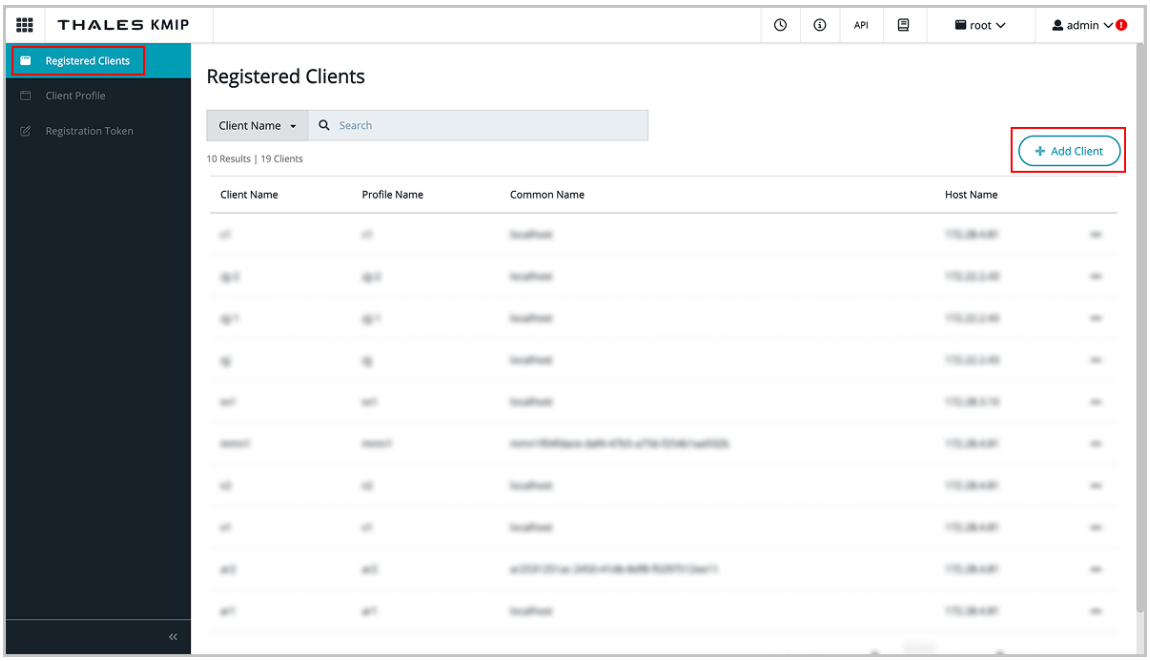
- In step 4 **Create Token**, copy the generated token.

Figure 10-16: Create a Token



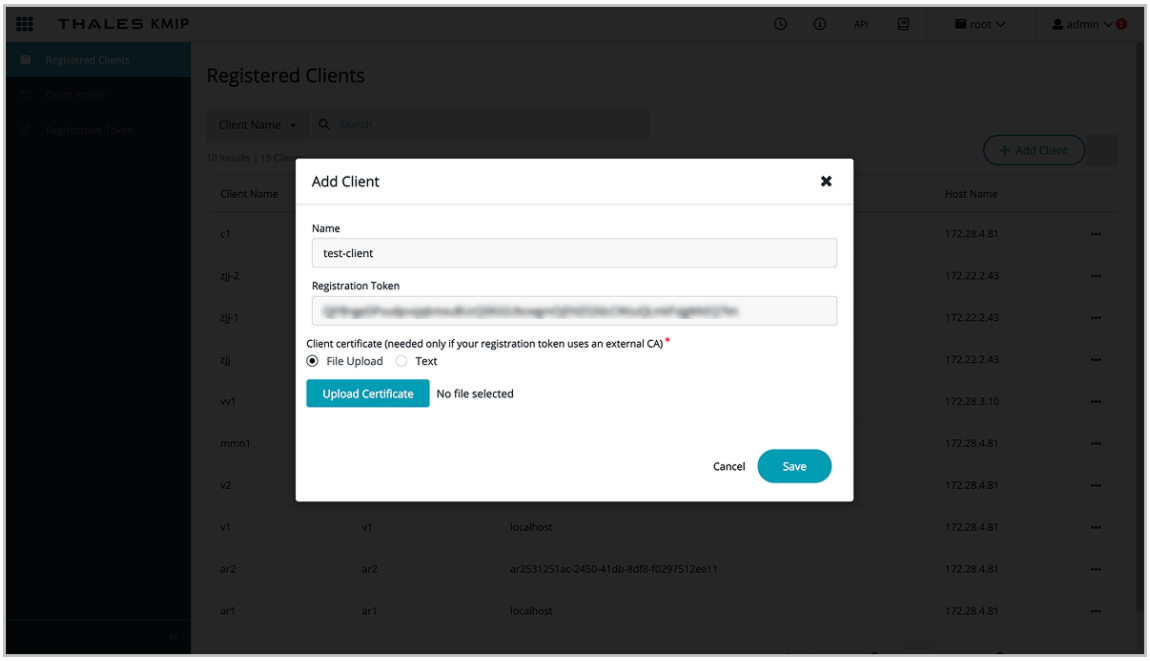
g) Click **Registered Clients** > **Add Client**.

Figure 10-17: Add a Client



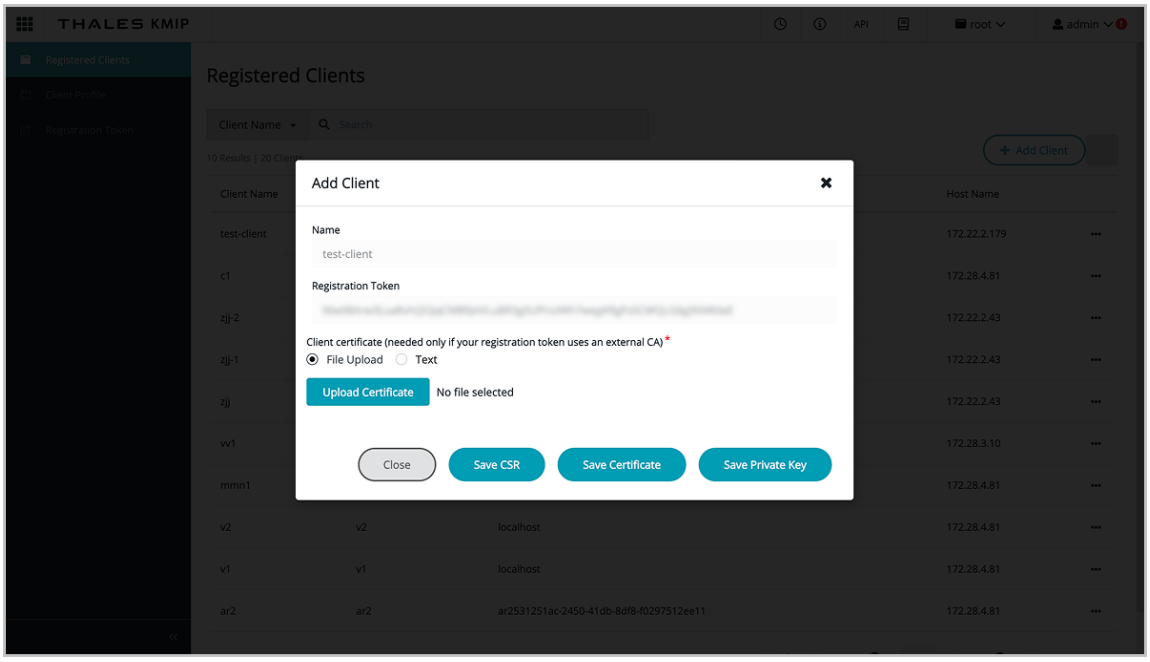
h) In the **Add Client** dialog, set the client name, paste the registration token, and then click **Save**.

Figure 10-18: Configure the Client



i) After saving successfully, click **Save Certificate** and **Save Private Key** to download the certificate and private key.

Figure 10-19: Download the Certificate and Private Key



j) Log in to NexaSphere.

k) Navigate to the **Key Providers** page.

l) Locate the Thales standard key provider you added, and click **Actions** > **Establish Trust** > **Make the KMS Trust NexaSphere**.

m) In the **Make the KMS Trust NexaSphere** dialog, select **KMS Certificate and Private Key** as the trust method.

n) Click **Next**.

o) Upload the KMS certificate and private key you downloaded from Thales.

p) Click **OK**.

10.6.1.5 Use the HashiCorp Key Provider

Context

1. *Add the HashiCorp key provider in NexaSphere*
2. *(Optional) Use the **KMS certificate and private key** method to make HashiCorp trust NexaSphere*
3. *(Optional) Use the **New Certificate Signing Request (CSR)** method to make HashiCorp trust NexaSphere*

Procedure

1. Add the HashiCorp key provider in NexaSphere.

   For detailed steps, see *Add a Standard Key Provider*.

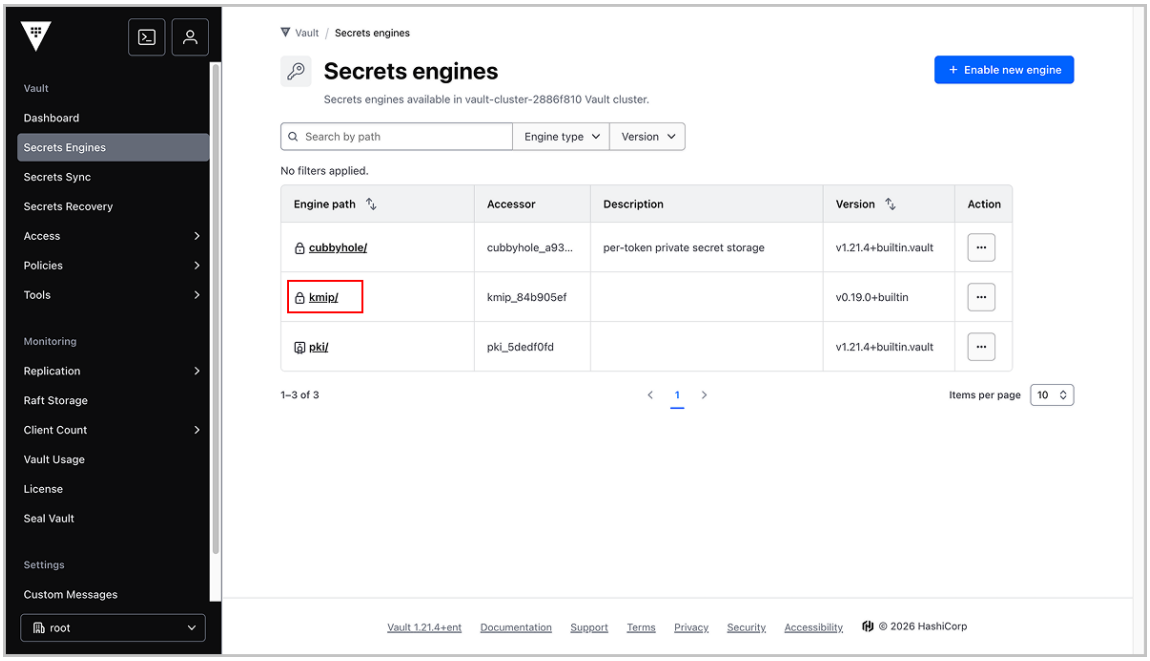
2. (Optional) Use the **KMS certificate and private key** method to make HashiCorp trust NexaSphere.

   a) Log in to HashiCorp.

   b) Click **Vault** > **Secrets Engines**.

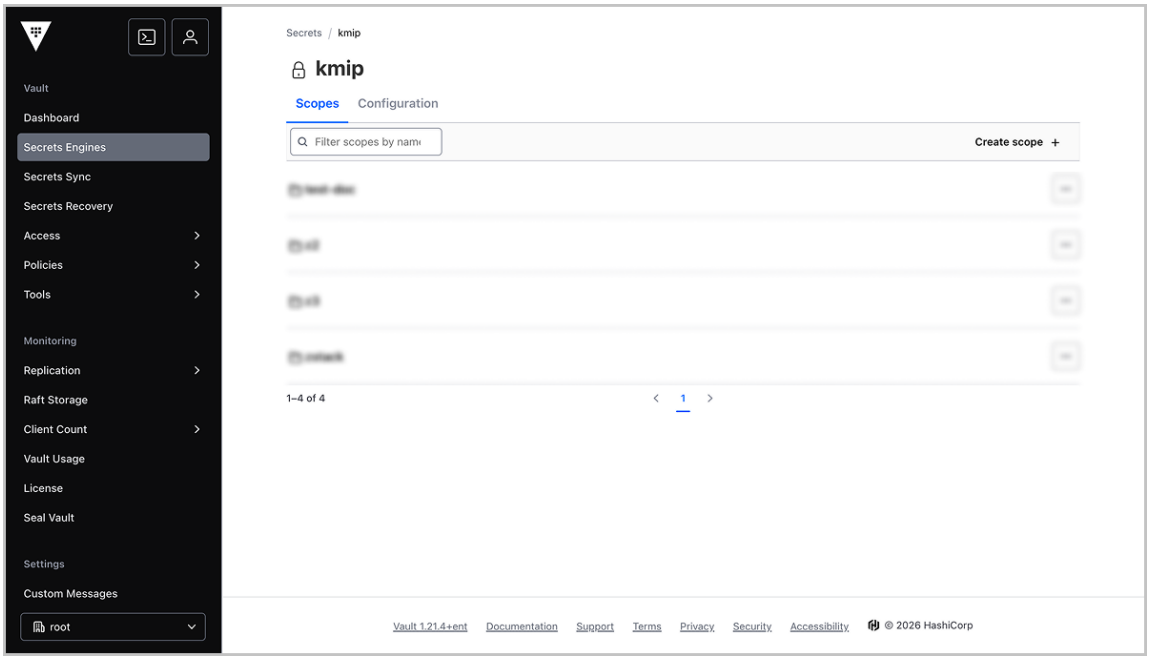
   c) On the **Secrets Engines** page, click **kmip** to access the KMIP module.

Figure 10-20: Access the KMIP Module



d) On the **kmip** page, click **Create scope** to create a scope or enter an existing one.

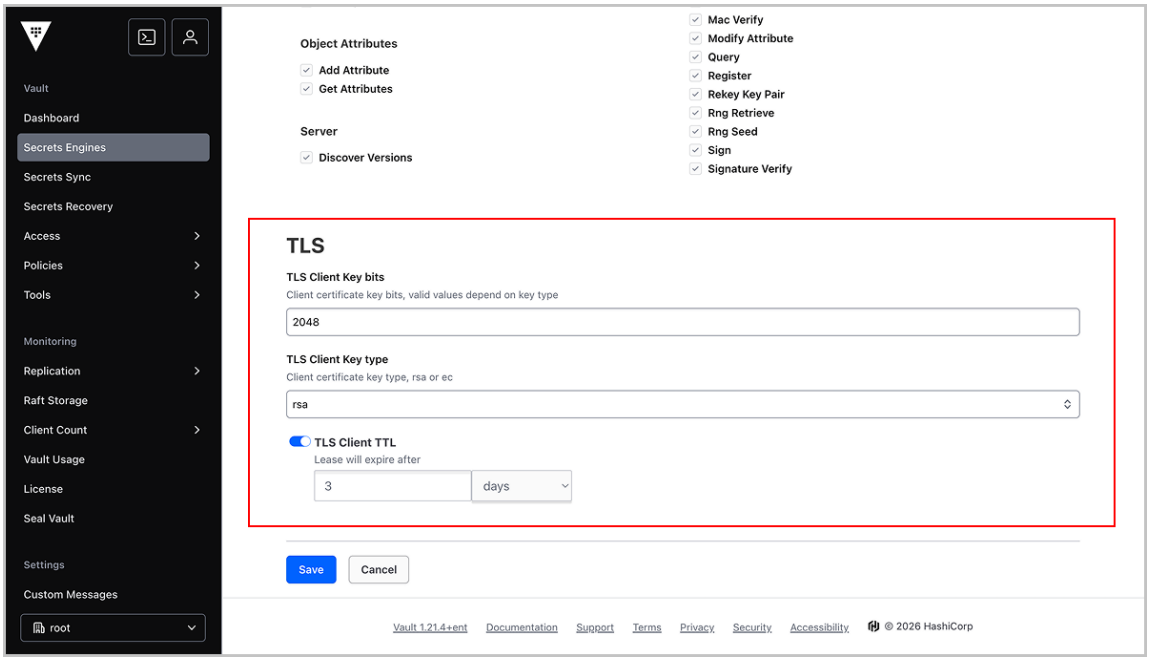
Figure 10-21: Create or Enter an Existing Scope



e) On the scope details page, click **Create role** to create a role.

f) On the **Create a Role** page, set the role name and complete the TLS configuration by referring to the example.

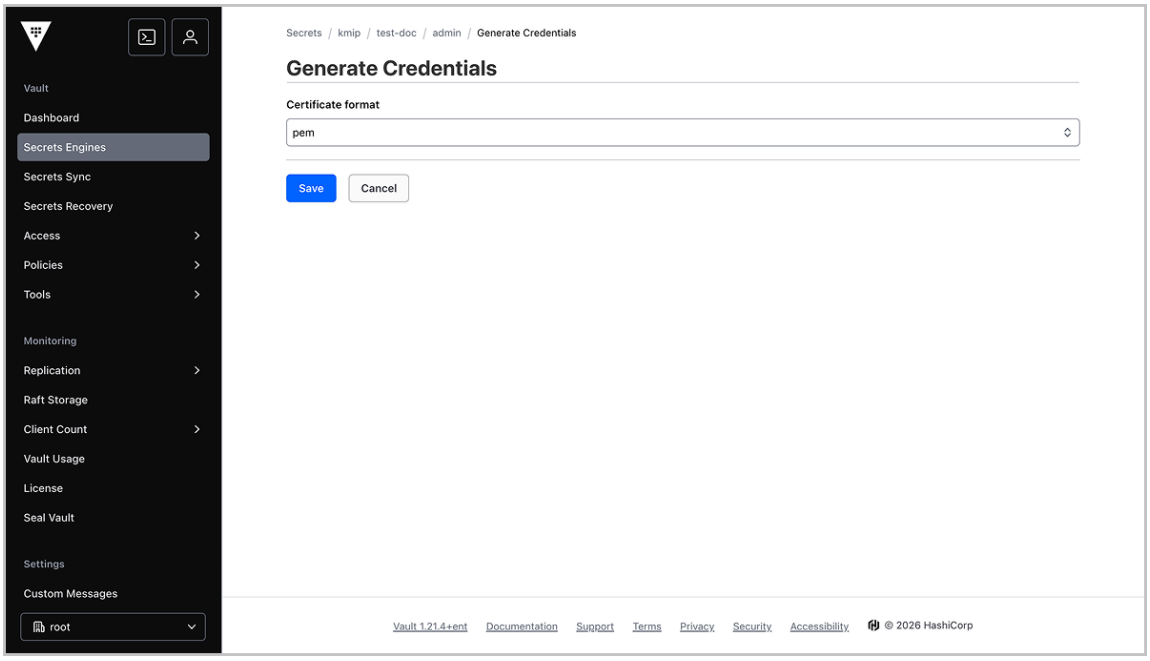
Figure 10-22: Create a Role



g) After you create the role, click the role to go to its details page, and then click **Generate credentials**.

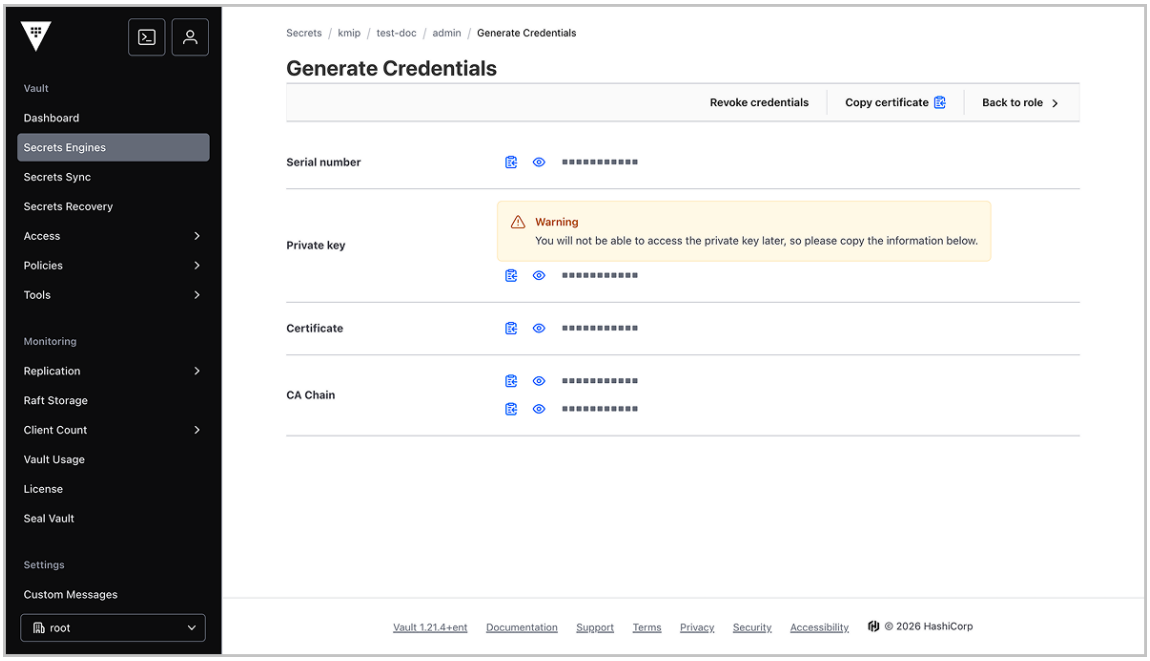
h) On the **Generate Credentials** page, select **pem** as the certificate format, and then click **Save**.

Figure 10-23: Generate a Certificate



i) Copy the certificate and private key.

Figure 10-24: Copy the Certificate and Private Key



j) Log in to NexaSphere.

k) Navigate to the **Key Providers** page.

l) Locate the HashiCorp standard key provider you added, and click **Actions** > **Establish Trust** > **Make the KMS Trust NexaSphere**.

m) In the **Make the KMS Trust NexaSphere** dialog, select **KMS Certificate and Private Key** as the trust method.

n) Click **Next**.

o) Paste the KMS certificate and private key you copied from HashiCorp.

p) Click **OK**.

3. (Optional) Use the **New Certificate Signing Request (CSR)** method to make HashiCorp trust NexaSphere.

a) Log in to the HashiCorp system.

b) Run `vault login` and enter the token.

c) Generate the CSR:

```
vault write -format=json kmip/scope/test/role/admin/credential/
sign csr="$(cat csr.pem)" | jq -r '.data.certificate' > client.crt
```

d) Log in to NexaSphere.

e) Navigate to the **Key Providers** page, and then click **Actions** > **Establish Trust** > **Upload Signed CSR Certificate**.

f) Upload the obtained `client.crt` certificate.

10.6.2 Configure and Manage a Native Key Provider

10.6.2.1 Add a Native Key Provider

NexaSphere provides a native key provider that allows you to enable encryption-related features without a third-party KMS.

Procedure

1. In the navigation pane, select **Inventory** > **VM and Host**.
2. Select the root node.
3. On the root node details page, click **Key Providers**.
4. On the **Key Providers** tab, click **Add Key Provider**.
5. In the **Add Key Provider** dialog, configure the following options:

Option	Description
Name	Enter a name for the key provider. **Note:** The name of each key provider serves as its identifier. Ensure the name is globally unique.
Description	Enter a brief description for the key provider.
Type	Select **Native Key Provider**.

6. Click **OK**.

   After the key provider is added successfully, its status shows as "Not Backed Up".

What's next

After you add a native key provider, you must back it up before use. For more information, see *Back Up the Native Key Provider*.

10.6.2.2 Back Up the Native Key Provider

The Native Key Provider is stored only in NexaSphere. If NexaSphere fails or must be rebuilt, its keys are lost and encrypted resources that depend on the provider become unavailable. Back

up the Native Key Provider so that you can restore the key service and regain access to the encrypted resources after rebuilding NexaSphere.

Procedure

1. In the navigation pane, select **Inventory** > **VM and Host**.
2. Select the root node.
3. On the root node details page, click **Key Providers**.
4. On the **Key Providers** tab, select the native key provider you want to back up.

   For key providers that are not yet backed up, the status shows as "Not Backed Up".
5. Click **Actions** > **Back Up**.
6. In the **Back Up Key Provider** dialog, select the backup mode as needed.

Option	Description
Without Password Protection	**Note:** Without password protection, the configuration data and the virtual machines encrypted with this key provider face potential security risks.
Password Protection	Enter and confirm the password. **Note:** Store the password in a safe place. The platform cannot access or recover the password you set. Make sure this password is securely saved, as it will be required to restore the Native Key Provider configurat ion in case of disaster. If you forget or lose the password: • You cannot restore the key provider. • You cannot access encrypted resources that depend on this key provider, such as vTPM-enabled VMs, encrypted VMs, and encrypted disks.

7. Click **OK**.

   The backup file is downloaded through your browser. Save the backup file in a secure location.

10.6.2.3 Restore the Native Key Provider

Prerequisites

- You have prepared the backup file of the native key provider.
- If you set password protection for the backup file when you backed up the key provider, prepare the password for the backup file.
- Ensure that no key provider with the same name exists on the platform.

Procedure

1. In the navigation pane, select **Inventory** > **VM and Host**.
2. Select the root node.
3. On the root node details page, click **Key Providers**.
4. On the **Key Providers** tab, click **Actions** > **Restore**.
5. In the **Restore Key Provider** dialog, upload the backup file.
6. (Optional) If the backup file is password-protected, enter the password.
7. Click **Next**.
8. Confirm the information of the key provider to restore, and then click **OK**.

10.6.2.4 Delete the Native Key Provider

Prerequisites

- If you want to keep the key provider for future use, back it up first.
- To preserve encrypted resources that depend on this key provider, perform a global **Rekey** operation before deletion. The operation uses a new KEK from the current default key provider to rewrap the DEKs of the encrypted resources.

Procedure

1. In the navigation pane, select **Inventory** > **VM and Host**.
2. Select the root node.
3. On the root node details page, click **Key Providers**.
4. On the **Key Providers** tab, select the native key provider you want to delete, and then click **Actions** > **Delete**.



Note:

- After deletion, resources that depend on this key provider, such as vTPM-enabled VMs, encrypted VMs, and encrypted disks, are locked and become unavailable.
- The key provider is permanently deleted from the platform and cannot be recovered. Proceed with caution.

5. Read the risk warning, enter the confirmation information, and click **OK**.

10.6.3 Rekey

To meet requirements such as periodic key rotation, suspected key compromise, or security compliance, perform a shallow rekey on encrypted resources across the platform. This operation uses a new key encryption key (KEK) from the current default key provider to rewrap the data encryption keys (DEKs). It does not replace the DEKs or re-encrypt the actual workload data.

Prerequisites

- Make sure that the current default key provider is connected and its key service is available.
- Make sure that the key providers and keys currently protecting encrypted resources across the platform are available.

Procedure

1. In the navigation pane, select **Inventory** > **VM and Host**.
2. Select the root node.
3. On the root node details page, click **Key Providers**.
4. On the **Key Providers** tab, click **Actions** > **Rekey**.
5. In the **Rekey** dialog, click **OK**.
   The platform completes a shallow rekey by using a new KEK from the current default key provider to rewrap the DEKs of encrypted resources, including VMs, disks, disk snapshots, vTPMs, and encrypted VMs and disks in the Recycle Bin. The actual workload data is not re-encrypted.

10.7 VM Encryption

NexaSphere uses key providers to protect VM runtime state and virtual disk data. VM encryption and disk encryption can be configured separately to protect an entire VM or individual disks.

VM encryption protects memory and GPU runtime state written to disk and protects data in transit during hot migration. Disk encryption protects data stored on virtual disks and in their snapshots. Encrypted resources depend on keys supplied by a key provider. If the key service becomes unavailable, the affected resources cannot operate normally.

10.7.1 VM Encryption and Disk Encryption

Protected Data

VM encryption Protects VM runtime state after it is written to disk. Hot migration of an encrypted VM always uses encrypted transport.

Disk encryption Protects data stored on virtual disks. Snapshots created from encrypted disks remain encrypted.

Encryption Configurations

VM encryption and disk encryption can be combined as follows:

Configuration	Description
VM encryption enabled	When you create a VM, encryption is enabled for all disks by default. You can disable encryption for individual disks as needed.
Disk encryption only	If VM encryption is disabled, you can still encrypt one or more individual disks.
Encryption disabled	The VM runtime state and disk data are not protected by the VM encryption feature.

10.7.2 Encryption Keys and How Encryption Works

Key Hierarchy

NexaSphere uses a hierarchy of keys to protect encrypted resources:

Data encryption key (DEK) Encrypts VM runtime state or virtual disk data.

Key encryption key (KEK) Is supplied by the key provider and protects the DEK, so the platform does not need to store an unprotected DEK.

Using and Protecting Key Providers

When you create an encrypted resource or change the encryption state of a resource, the platform obtains keys from the default key provider. After an encrypted resource is created, it continues to depend on the key provider and key that protect its DEK.

Keep keys available based on the key provider type:

Standard Key Provider	Maintain connectivity between the platform and the key provider, and keep the key service and associated keys available.
Native Key Provider	Its keys are stored with the platform. Before using a Native Key Provider to create encrypted resources, back it up and store the backup file in a secure location. If the platform fails or must be rebuilt, use the backup to restore the key service. For instructions, see *Back Up the Native Key Provider*.

If the key provider connection fails, the key service becomes unavailable, or keys are lost, the encrypted resources that depend on the key provider are locked and become unavailable.

Rekeying Encrypted Resources

Changing the default key provider does not automatically update the KEKs used by existing encrypted resources. Administrators can perform a shallow rekey to rotate the KEKs that protect their DEKs or, after changing the default key provider, to protect existing DEKs with new KEKs from the current default key provider. A shallow rekey does not replace the DEKs or re-encrypt the actual data in VMs or disks.

A global rekey covers VMs, disks, disk snapshots, vTPMs, and encrypted VMs and disks in the Recycle Bin.

10.7.3 Encryption Requirements and Limitations

Requirements

Item	Requirement
License	Make sure the platform has a valid Advanced Edition license.
Key provider	A default key provider must be configured, connected, and available.

Configuration and Storage Limitations

Item	Limitation
Changing encryption	Power off an existing VM before enabling or disabling VM encryption or disk encryption. Enabling VM encryption automatically enables encryption for existing disks. Disabling VM encryption does not automatically disable disk encryption; you can change the encryption state of each disk separately.

Item	Limitation
	Converting existing data usually takes longer than enabling encryption when the VM is created.
Shared disks	Disk encryption and disk sharing are mutually exclusive. A shared disk cannot be encrypted, and an encrypted disk cannot be configured as a shared disk.
New disks and disk images	When you create a disk or add one from a disk image, disk encryption follows the VM encryption state by default and can be changed as needed.
Existing disks	When you attach an existing disk to a VM, the disk retains its encryption state, which cannot be changed during attachment. The VM encryption state does not change the encryption state of the existing disk.
RDM disks	RDM disks do not support disk encryption and are not automatically encrypted when VM encryption is enabled.
Data storage	An encrypted disk can be created or migrated only on data storage that supports disk encryption. ZHPS distributed storage does not support disk encryption.
Data cleanup	After you change the encryption state of an existing VM or disk, the source data from before the conversion is retained on the **Data Cleanup** tab of the data storage. Verify that the converted data is usable, and then remove the source data.

Interoperability with Other Features

Operation	Scenario	Requirement or Limitation
Cloning	Clone a VM	If VM encryption is enabled on the source VM, it remains enabled on the cloned VM. Encrypted source disks remain encrypted. Unencrypted source disks remain unencrypted by default and can be encrypted as needed. If VM encryption is enabled on the cloned VM, encryption is enabled by default for its unencrypted disks.
Snapshots	Create a snapshot	A snapshot inherits the encryption state of each disk. Snapshot data created for an encrypted disk remains

Operation	Scenario	Requirement or Limitation
		encrypted. A snapshot does not record the encryption state of VM runtime data.
	Create a VM from a snapshot	The VM encryption option follows the current VM encryption state of the source VM. Each disk inherits the encryption state of its corresponding snapshot disk. An encrypted snapshot disk remains encrypted and cannot be decrypted during creation. Encryption can be enabled for an unencrypted snapshot disk as needed.
Backup	Create backup data	When a full or incremental backup is created for an encrypted VM or disk, the corresponding backup data remains encrypted. The default key provider must be available when encrypted backup data is created.
	Synchronize backup data	Backup data for a VM with VM or disk encryption enabled cannot be synchronized to remote backup storage.
Migration	Hot-migrate a VM with VM encryption enabled	Migration traffic encryption is enabled automatically and cannot be disabled, regardless of the disk encryption state.
	Hot-migrate a VM without VM encryption	Migration traffic encryption is disabled by default and can be enabled as needed. Disk encryption alone does not require migration traffic encryption.
	Perform a cold migration that transfers disk data	When the migration interface provides the migration traffic encryption option, it is enabled automatically and cannot be disabled if VM encryption is enabled. Otherwise, it is disabled by default and can be enabled as needed.
	Perform a cold migration that does not transfer VM runtime state or disk data	Migration traffic encryption does not need to be configured. For example, changing only the host of a shut-down VM on shared storage does not transfer disk data.
	Change the data storage of an encrypted disk	An encrypted disk can be migrated only to data storage that supports encryption.
	After migration	The VM and each disk retain their respective encryption states from before migration.
	Perform a batch migration	If any selected VM has VM encryption enabled, migration encryption is enabled for the entire batch and cannot be disabled.

Operation	Scenario	Requirement or Limitation
High Availability	Recover a VM through HA	The VM and its disks retain their encryption states after HA. The key provider and keys that protect the resources must remain available.
Templates and Export	Convert to a template	These operations are not supported for VMs that have VM encryption or disk encryption enabled.
	Clone to a template	
	Export as an OVA file	

10.8 VM Trusted Security

VM trusted security is used to enhance the security of the VM operating system runtime environment. By configuring key providers, vTPM, UEFI boot, and Secure Boot, VMs can obtain trusted platform module capabilities and stronger boot chain protection.

Before you configure vTPM for a VM, make sure an available key provider is configured on the platform and that the VM runtime environment and operating system meet the requirements for TPM 2.0, UEFI, and related features.

10.8.1 Secure VMs with vTPM

When you create a new virtual machine or add a virtual Trusted Platform Module (vTPM) to an existing virtual machine, you provide enhanced security for the virtual machine operating system.

Prerequisites

- The platform has an available key provider. If you use the native key provider, back it up first.
- The host on which the VM runs must meet the vTPM support requirements for kernel, QEMU, and libvirt versions. The following table lists the required component versions for each architecture and host operating system version:

Architecture	Host OS version	Kernel version	QEMU version	libvirt version
x86	H84r	`4.18.0-553.69.1.13.gc220c6303994.el8.x86_64`	`qemu-kvm-6.2.0-235.g51749aa16b.el8`	`libvirt-8.0.0-90.g4f8dd1cb01.el8.x86_64`
	KY10 SP3	`4.19.90-52.48.v2207.ky10.x86_64`	`qemu-kvm-6.2.0-235.g51749aa16b.ky10`	`libvirt-8.0.0-90.g4f8dd1cb01.ky10.x86_64`

Architecture	Host OS version	Kernel version	QEMU version	libvirt version
	KY10 SP3. 2403	`4.19.90-89.25. v2401.ky10.x86_64`	`qemu-kvm-6.2.0-232.g09252161d1.ky10`	`libvirt-8.0.0-90.g4f8dd1cb01.ky10.x86_64`
ARM	KY10 SP3	`4.19.90-52.48. v2207.ky10.aarch64`	`qemu-6.2.0-906.g271454a05e.ky10`	`libvirt-6.2.0-415.gf2c25be909.ky10.aarch64`
	KY10 SP3. 2403	`4.19.90-89.25. v2401.ky10.aarch64`	`qemu-6.2.0-1042.g7eecd245dd.ky10`	`libvirt-6.2.0-463.g2ee090fe60.ky10.aarch64`
	H22e	`5.10.0-136.12.0.86 .oe2203sp1.aarch64`	`qemu-6.2.0-906.g271454a05e`	`libvirt-6.2.0-415.gf2c25be909.aarch64`

- The virtual machine boot mode must be set to UEFI.
- The virtual machine operating system must support the TPM 2.0 specification. Operating systems that do not support TPM 2.0 (including but not limited to Windows Server 2012 and earlier, CentOS 6 / RHEL 6, Ubuntu 14.04 and earlier) are not supported.
- Before you add a TPM to an existing virtual machine, power off the virtual machine.

Procedure

1. In the navigation pane, select **Inventory** > **VM and Host**.
2. In the resource tree, right-click the target cluster, host, or image, and then click **New Virtual Machine**.
3. In the **Select VM Creation Method** dialog, select **New VM**, and then click **Next**.
4. In the **New Virtual Machine** dialog, complete the relevant basic configuration. For more information about virtual machine parameters, see *Create a Virtual Machine*.
5. In the **Advanced Settings** section, click **Boot Options** and select **UEFI** as the virtual machine BIOS mode.
6. In the **Hardware Info** section, click **Add Hardware** > **TPM**.
   After the TPM is added successfully, click TPM in the hardware list. You can view the TPM status and specification in the hardware configurations pane on the right.
7. After you confirm the configuration details are correct, click **OK**.

11 Data Protection and Recovery

11.1 Overview

The backup management module provides a set of capabilities for data protection and disaster recovery.

This book describes the main backup management features and provides guidance for common data protection and recovery operations.

Feature Overview

Centralized Management Interface

NexaSphere provides a centralized UI for managing protected resources, backup data, backup plans, and backup storage.

Plan-Based Backup

The backup service in NexaSphere is based on periodic scheduled backups. With backup policies, retention policies, and backup QoS, you can create backup plans that meet your business requirements and apply them to virtual machines or the platform database.

Multiple Backup Policies

NexaSphere provides multiple backup policies for different backup objects. For virtual machines, NexaSphere supports default incremental backup and custom incremental backup policies. For the platform database, you can customize the backup cycle, execution time, and start time.

Flexible Backup Scenarios

NexaSphere supports local and remote data protection. Local backup storage supports seamless failover for business continuity, and backup data can be synchronized from local backup storage to remote backup storage.

Backup Activity Monitoring

NexaSphere provides detailed information about backup activities, including backup plan execution status and backup data capacity usage.

11.2 Prerequisites

Before you use the backup feature in NexaSphere, make sure that the license meets the requirements. Backup licenses take effect according to the following rules:

- A new installation includes a three-month trial license for backing up up to six VMs.

- In an upgraded environment, the backup license remains unchanged. If the original preset license was still in use before the upgrade, its quota and validity period remain unchanged after the upgrade.
- After a license update, the backup feature uses the quota and validity period provided by the new license.

After the number of protected VMs reaches the trial quota, purchase a module license to continue using the backup feature. NexaSphere provides backup licenses by VM count or with an unlimited quota. The purchased quota replaces the trial quota. For example, if the trial quota is 6 and the purchased quota is 50, the effective quota is 50.

11.3 Add Backup Storage

To use the backup feature, deploy backup storage in NexaSphere to store backup data. Depending on your backup scenario, you can add local backup storage or remote backup storage.

11.3.1 Add Local Backup Storage

Add local backup storage to store scheduled backup data for local virtual machines or the platform database. Local backup storage supports seamless failover to help ensure business continuity.

If local data is accidentally deleted or corrupted, you can restore backup data from local backup storage to the local environment.

If a disaster occurs in the local data center, you can use local backup storage to rebuild the data center and recover services.

You can add local backup storage in any of the following ways:

- Reuse an existing image storage resource to store backup data. For more information, see *Reuse Image Storage*.
- Reuse an existing host resource and use the free capacity on the host to store backup data. For more information, see *Reuse Host*.
- Add a dedicated backup node and use free disks or a local directory on the node to store backup data. For more information, see *Dedicated Backup Storage*.

You can configure the number of local backup storage resources based on your business requirements.

11.3.1.1 Reuse Image Storage

Reuse an existing image storage resource to store local backup data.

Prerequisites

- An image storage resource already exists in the platform and is a connected and enabled standalone image repository.
- The backup network has been planned.

Procedure

1. In the navigation pane, choose **Data Protection** > **Backup Storage**.
2. On the **Backup Storage** page, click **Add Backup Storage**.
3. In the **Select Backup Storage Addition Type** dialog, select **Reuse Image Storage**, and then click **Next**.
4. In the **Add Local Backup Storage** dialog, configure the following options:

   Basic Information

Option	Description
Data Center	Displays the data center where the backup storage is located.
Addition Method	Displays **Reuse Image Storage**.

   Configuration Information

Option	Description
Image Storage	Select an existing image storage to use as local backup storage.
Backup Storage Path	Automatically obtains the path of the selected image storage.
Backup Network	The network dedicated to backup traffic. Enter the backup network CIDR.
Backup Data	Choose whether to scan existing backup data.

5. Review the configuration and click **OK**.

11.3.1.2 Reuse Host

Reuse an existing host resource and use the free capacity on the host to store local backup data.

Prerequisites

- A host resource already exists in the platform and is connected and enabled.
- The backup network has been planned.

Procedure

1. In the navigation pane, choose **Data Protection** > **Backup Storage**.
2. On the **Backup Storage** page, click **Add Backup Storage**.
3. In the **Select Backup Storage Addition Type** dialog, select **Reuse Host**, and then click **Next**.
4. In the **Add Local Backup Storage** dialog, configure the following options:

   Basic Information

Option	Description
Name	Name of the local backup storage.
Description	Enter a brief description for the local backup storage.
Data Center	Displays the data center where the backup storage is located.
Addition Method	Displays **Reuse Host**.

   Configuration Information

Option	Description
Host	Select an existing host to use as local backup storage.
Backup Storage Path	Enter the mount path on the local backup storage. **Note:** Do not use system directories such as `/`, `/dev/`, `/proc/`, `/sys/`, `/usr/bin`, or `/bin`. Otherwise, the host might become abnormal.
Backup Network	The network dedicated to backup traffic. Enter the backup network CIDR.
Backup Data	Choose whether to scan existing backup data.

5. Review the configuration and click **OK**.

11.3.1.3 Dedicated Backup Storage

Add a dedicated backup node and use free disks or a local directory on the node to store backup data.

Procedure

1. In the navigation pane, choose **Data Protection** > **Backup Storage**.
2. On the **Backup Storage** page, click **Add Backup Storage**.
3. In the **Select Backup Storage Addition Type** dialog, select **Dedicated Backup Storage**, and then click **Next**.
4. In the **Add Local Backup Storage** dialog, complete the basic configuration and disk configuration.
5. For **Basic Configuration**, configure the following options:

   Basic Information

Option	Description
Name	Name of the local backup storage.
Description	Enter a brief description for the local backup storage.
Data Center	Displays the data center where the backup storage is located.
Addition Method	Displays **Dedicated Backup Storage**.

   Configuration Information

Option	Description
Backup Storage IP	Enter the IP address of the dedicated backup storage.
SSH Port	Default is 22.
Username	Default is the root user.
Password	The password of the specified user.

6. Review the basic configuration and click **Next**.
7. For **Disk Configuration**, configure the following options:

Option	Description
Storage Method	Supports **Free Disk** and **Local Directory**.

Option	Description
Free Disk	Set this parameter when **Free Disk** is selected as the storage method. **Note:** Configuring a free disk formats the selected disk and permanently clears all partitions, file systems, and data from the disk.
Backup Storage Path	Enter the mount path on the dedicated backup storage. **Note:** Do not use system directories such as `/`, `/dev/`, `/proc/`, `/sys/`, `/usr/bin`, or `/bin`. Otherwise, the backup storage might become abnormal.
Backup Network	The network dedicated to backup traffic. Enter the backup network CIDR.
Backup Data	Set this parameter when **Local Directory** is selected as the storage method. Choose whether to scan existing backup data.

8. Review the configuration and click **OK**.

11.3.2 Add Remote Backup Storage

Add remote backup storage to store scheduled backup data for local virtual machines or the platform database. Backup data is synchronized from local backup storage to remote backup storage.

If local data is accidentally deleted or corrupted, you can restore backup data from remote backup storage to the local environment.

If a disaster occurs in the local data center, you can use remote backup storage to rebuild the data center and recover services.

You can add remote backup storage by adding a dedicated backup node. For more information, see *Dedicated Backup Storage*.

You can add only one remote backup storage.

11.3.2.1 Dedicated Backup Storage

Add remote backup storage by adding a dedicated backup node.

Procedure

1. In the navigation pane, choose **Data Protection** > **Backup Storage**.
2. On the **Backup Storage** page, click **Add Backup Storage**.
3. In the **Select Backup Storage Addition Type** dialog, select **Dedicated Backup Storage**, and then click **Next**.
4. In the **Add Remote Backup Storage** dialog, complete the basic configuration and disk configuration.
5. For **Basic Configuration**, configure the following options:

   Basic Information

Option	Description
Name	Name of the remote backup storage.
Description	Enter a brief description for the remote backup storage.
Data Center	Displays the data center where the backup storage is located.
Addition Method	Displays **Dedicated Backup Storage**.

   Configuration Information

Option	Description
Backup Storage IP	Enter the IP address of the dedicated backup storage.
SSH Port	Default is 22.
Username	Default is the root user.
Password	The password of the specified user.

6. Review the basic configuration and click **Next**.
7. For **Disk Configuration**, configure the following options:

Option	Description
Storage Method	Supports **Free Disk** and **Local Directory**.

Option	Description
Free Disk	Set this parameter when **Free Disk** is selected as the storage method. **Note:** Configuring a free disk formats the selected disk and permanently clears all partitions, file systems, and data from the disk.
Backup Storage Path	Enter the mount path on the dedicated backup storage. **Note:** Do not use system directories such as `/`, `/dev/`, `/proc/`, `/sys/`, `/usr/bin`, or `/bin`. Otherwise, the backup storage might become abnormal.
Backup Network	The network dedicated to backup traffic. Enter the backup network CIDR.
Backup Data	Set this parameter when **Local Directory** is selected as the storage method. Choose whether to scan existing backup data.

8. Review the configuration and click **OK**.

11.4 Virtual Machine Backup

The backup service in NexaSphere is based on periodic scheduled backups. To perform backups, configure a backup plan first. For more information, see *Create a Backup Plan*.

During backup plan execution, the backup mode, backup chain length, backup data retention policy, and backup QoS depend on the configured backup policy. For more information, see *Backup Policy*.

When a full or incremental backup is created for a VM with VM encryption or disk encryption enabled, the corresponding backup data remains encrypted. Before creating encrypted backup data, make sure that the default key provider is connected and available.

After backups are generated, you can use backup data to recover virtual machines. For more information, see *Data Recovery*.

11.4.1 Backup Policy

Backup Policy

Backup Mode

NexaSphere provides two backup modes:

- Default incremental backup

  You only need to customize the incremental backup policy. The full backup policy is configured by the system. After 63 incremental backups are performed, the system automatically performs one full backup, and the maximum backup chain length is 64.

- Custom incremental backup

  You need to customize both the incremental backup policy and the full backup policy. By customizing the full backup policy, you can shorten the full backup interval. More full backups shorten the backup chain and reduce the amount of backup data to merge during recovery.

Backup Cycle

The backup cycle determines how frequently backups are performed.

- For incremental backup policies, NexaSphere provides multiple options, including monthly, weekly, daily, hourly, and minutely backups.
- For full backup policies, NexaSphere supports monthly and weekly backups.

Execution Time

You can configure a more granular backup execution time down to the minute. If the backup cycle of an incremental backup policy is monthly, weekly, or daily, you can set the incremental backup execution time. If the backup cycle of a full backup policy is monthly or weekly, you can set the full backup execution time.

- Monthly backup: If the execution time is set to **00:00 on the 1st day of every month**, the backup starts at 00:00:00 on the first day of each month.
- Weekly backup: If the execution time is set to **Sunday Tuesday 00:00**, the backup starts at 00:00:00 every Sunday and Tuesday.
- Daily backup: If the execution time is set to **00:30**, the backup starts at 00:30:00 every day.

Start Time

The backup plan starts at the start time and performs backups based on the configured backup cycle and execution time.

Retention Policy

Each successful backup generates backup data that can be used to restore VM data to an earlier point in time. To control the amount of backup data, specify a backup data retention policy. Based on the backup scenario, NexaSphere can set retention policies separately for local backup data and remote backup data.

- Local retention policy:
  - Retain by quantity: Retain at least the latest local incremental backup data and the latest local full backup data.
  - Retain by time: Retain local backup data for at least the latest 1 day.
- Remote retention policy:
  - Permanent retention: Remote backup data is not automatically cleaned up.
  - Retain by quantity: Retain at least the latest remote incremental backup data and the latest remote full backup data.
  - Retain by time: Retain remote backup data for at least the latest 1 day.

 **Note:**

Backup data outside the retention policy is removed only from the records. Data security is not affected.

11.4.2 Create a Backup Plan

Prerequisites

- Make sure the platform has sufficient authorized backup quota.
- Make sure the virtual machines to be backed up are in the running or paused state and are not associated with any other backup plans.
- If the data storage where a virtual machine resides is ZHPS distributed storage, the virtual machine cannot be associated with a backup plan.

Procedure

1. In the navigation pane, choose **Data Protection** > **Backup Plan**.
2. On the **Backup Plan** page, click **New Backup Plan**.
3. In the **New Backup Plan** dialog, complete the basic configuration and backup configuration.
4. For **Basic Configuration**, configure the following options:

Basic Information

Option	Description
Name	Name of the backup plan.
Description	Enter a brief description for the backup plan.
Data Center	Select a data center.

Backup Storage

Option	Description
Local Backup Storage	Add local backup storage. You can specify up to two backup storage resources for seamless failover.
Sync to Remote Backup Storage	Disabled by default. Choose whether to synchronize local backup data to remote backup storage. **Note:** If this option is enabled, add remote backup storage to the platform in advance. For more information, see *Add Remote Backup Storage*.
Remote Backup Storage	After **Sync to Remote Backup Storage** is enabled, displays the name and available capacity of the remote backup storage.

Backup Object

Option	Description
Backup Object Type	Select **Virtual Machine**.
Virtual Machine	Add virtual machines to be backed up and set the backup priority as needed. **Note:** • The number of selected virtual machines must not exceed the authorized backup quota.

Option	Description
	• If a selected virtual machine already has backup data, creating a new backup plan does not consume additional authorized backup quota. • By default, an entire VM is backed up. If the VM has shared disks or RDM disks attached, the shared disks and RDM disks are not backed up.

5. Review the basic configuration and click **Next**.
6. For **Backup Configuration**, configure the following options:

   Backup Policy

Option	Description
Backup Mode	Select default incremental backup or custom incremental backup as needed.
Incremental Backup Policy	Customize the incremental backup policy, including the backup cycle and execution time.
Full Backup Policy	When the backup mode is custom incremental backup, customize the full backup policy, including the backup cycle and execution time.
Maximum Backup Chain Length	Displays the backup chain length under the current backup policy. The maximum length cannot exceed 64.
Start Time	Set the start time of the backup plan.
Disk Read Speed	Set the disk read speed limit. If this parameter is left empty, the speed is unlimited by default. Unit: MB/s or GB/s. Value range: 1 MB/s to 100 GB/s. **Note:** Configure QoS based on the physical network bandwidth and the bandwidth consumed by concurrent backups.

   For more information about backup policies, see *Backup Policy*.

Retention Policy

Option	Description
Local Retention Policy	Set the local backup retention policy by quantity or by time.
Remote Retention Policy	Set this parameter after **Sync to Remote Backup Storage** is enabled. Set the remote backup retention policy to permanent retention, retention by quantity, or retention by time.

For more information about retention policies, see *Backup Policy*.

7. Review the configuration and click **OK**.

What's next

After backups are generated, you can use backup data to recover virtual machines. For more information, see *Data Recovery*.

11.4.3 View Backup Plans and Data

During backup plan execution, you can view the execution status of backup plans to monitor backup activity in real time. For more information, see *View Backup Plan Executions*.

During backup plan execution, backup data is generated. You can view backup data for protected resources. For more information, see *View VM Backup Data*.

11.4.3.1 View Backup Plan Executions

View the Latest Backup Result

Procedure:

1. In the navigation pane, choose **Data Protection** > **Backup Plan**.
2. In the list, view the latest backup result from **Latest Backup Job Result**.

Check Backup Execution Status of Protected Resources

If you start a large backup plan that involves multiple virtual machines, you can view the backup result of each virtual machine to identify which virtual machines succeeded, failed, or are waiting to be backed up.

Procedure:

1. In the navigation pane, choose **Data Protection** > **Backup Plan**.

2. On the **Backup Plan** page, click the backup plan name to go to the details page.
3. On the backup plan details page, click **Backup Resources**.
4. On the **Backup Resources** tab, view the backup execution status of all protected resources in the backup plan.

View Backup Job Details in a Backup Plan

You can view details of each backup job, including the backup object, backup mode, number of backup resources, total backup capacity, backup job result, backup duration, start and completion time, and backup details.

Procedure:

1. In the navigation pane, choose **Data Protection** > **Backup Plan**.
2. On the **Backup Plan** page, click the virtual machine backup plan name to go to the details page.
3. On the backup plan details page, click **Backup Job**.
4. On the **Backup Job** tab, select the backup job that you want to view, and then click **View Backup Details**.
5. In the **Backup Job Details** dialog, view details of the backup job.

11.4.3.2 View VM Backup Data

Backup plan execution generates a series of backup data records, which are stored in the specified backup storage. NexaSphere provides a unified entry for viewing backup data of protected resources so that important protected workloads are not missed.

Procedure

1. In the navigation pane, choose **Data Protection** > **Protected Resources**.
2. On the **Protected Resources** page, select the **Virtual Machine** tab.
3. View VM backup data.

   The VM backup data page displays backup data by virtual machine in a tree folder structure. Under each VM folder, all backup data for that VM is listed by creation time, with the latest backup data displayed first.

   - After selecting a VM, you can view the VM power state, total backup data capacity, associated backup plan, UUID, and all backup data of the VM.
   - Backup data that contains an encrypted VM or encrypted disk is identified by an encryption icon next to its name.

- After selecting backup data of the VM, you can view basic information and backup information about the backup data, including the source VM, backup capacity, backup type , backup storage, synchronization status to remote backup storage, owner, UUID, creation time, and detailed backup information.

11.4.4 Manual Backup

11.4.4.1 Run a Backup Now

Use manual backup if you want to create an additional backup restore point for one or more virtual machines in a backup plan, do not want to create a new backup plan or modify an existing one, or have associated a backup plan with the VMs to be protected but backup data has not yet been generated based on the scheduled time.

Procedure

1. In the navigation pane, choose **Data Protection** > **Backup Plan**.
2. On the **Backup Plan** page, select the backup plan for which you want to manually run a backup, and then click **Actions** > **Backup Now**.
3. In the **Backup Now** dialog, choose whether to enable **Full Backup**.

   Note:
   If enabled, a full backup is manually performed for the backup objects. If disabled, the backup mode is the same as that of the next scheduled backup.
4. Review the information and click **OK**.

   If the backup plan contains multiple virtual machines and you want to generate backup data for only one specific virtual machine, go to the **Backup Resources** tab on the backup plan details page, select the virtual machines that you do not want to back up, and then click **Actions** > **Backup Job State** > **Disable**. After they are disabled, manually run the backup plan.

11.4.4.2 Create an On-Demand Backup

In addition to backing up a VM by adding it to a backup plan, you can create an on-demand backup for the VM.

Prerequisites

- If the data storage where the VM resides is ZHPS distributed storage, backup creation is not supported.

- If the VM has an RDM disk attached, you cannot create an on-demand backup. To create a backup, detach the RDM disk first.
- If the VM has a shared disk attached, the shared disk is not backed up.

Procedure

1. In the navigation pane, choose **Inventory** > **Hosts and VMs**.
2. In the left navigation tree, right-click the VM and choose **Backup** > **Create Backup**.
3. In the **Create Backup** dialog, enter a backup name, select a backup type, specify local backup storage, and choose whether to synchronize the backup to remote backup storage.

   Note:
   If VM or disk encryption is enabled, **Sync to Remote Backup Storage** is unavailable.
4. Review the information and click **OK**.

What's next

After the backup is created, go to the **Protected Resources** page to view the backup data of the VM.

11.5 Platform Database Backup

The backup service in NexaSphere is based on periodic scheduled backups. To perform backups, configure a backup plan first. For more information, see *Create a Backup Plan*.

During backup plan execution, the backup cycle, execution time, start time, and backup data retention policy depend on the configured backup policy. For more information, see *Backup Policy*.

After backups are generated, you can use backup data to restore the platform database. For more information, see *Data Recovery*.

11.5.1 Backup Policy

Backup Policy

Backup Cycle

The backup cycle determines how frequently backups are performed. For platform database backups, NexaSphere provides three options: weekly, daily, and hourly backups.

Execution Time

If the backup cycle is weekly or daily, you can configure a more granular backup execution time down to the minute.

- Weekly backup: If the execution time is set to **Sunday Tuesday 00:00**, the backup starts at 00:00:00 every Sunday and Tuesday.
- Daily backup: If the execution time is set to **00:30**, the backup starts at 00:30:00 every day.

Start Time

The backup plan starts at the start time and performs backups based on the configured backup cycle and execution time.

Retention Policy

Each successful backup generates backup data that can be used to restore the platform database to an earlier point in time. To control the amount of backup data, specify a backup data retention policy. Based on the backup scenario, NexaSphere can set retention policies separately for local backup data and remote backup data.

- Local retention policy:
  - Retain by quantity: Retain at least the latest local backup data.
  - Retain by time: Retain local backup data for at least the latest 1 day.
- Remote retention policy:
  - Permanent retention: Remote backup data is not automatically cleaned up.
  - Retain by quantity: Retain at least the latest remote backup data.
  - Retain by time: Retain remote backup data for at least the latest 1 day.

Note:
Backup data outside the retention policy is removed only from the records. Data security is not affected.

11.5.2 Create a Backup Plan

Prerequisites

- Make sure the platform has sufficient authorized backup quota.
- NexaSphere supports only one platform database backup plan. Make sure no platform database backup plan already exists.

Procedure

1. In the navigation pane, choose **Data Protection** > **Backup Plan**.
2. On the **Backup Plan** page, click **New Backup Plan**.

3. In the **New Backup Plan** dialog, complete the basic configuration and backup configuration.
4. For **Basic Configuration**, configure the following options:

Basic Information

Option	Description
Name	Name of the backup plan.
Description	Description of the backup plan.
Data Center	Select a data center.

Backup Storage

Option	Description
Local Backup Storage	Add local backup storage. You can specify up to two backup storage resources for seamless failover.
Sync to Remote Backup Storage	Disabled by default. Choose whether to synchronize local backup data to remote backup storage. **Note:** If this option is enabled, add remote backup storage to the platform in advance. For more information, see *Add Remote Backup Storage*.
Remote Backup Storage	After **Sync to Remote Backup Storage** is enabled, displays the name and available capacity of the remote backup storage.

Backup Object

Option	Description
Backup Object Type	Select **Platform Database**. **Note:** If a platform database backup plan already exists, the Platform Database option is not displayed.

5. Review the basic configuration and click **Next**.

6. For **Backup Configuration**, configure the following options:

   Backup Policy

Option	Description
Backup Cycle	Select a backup cycle as needed, including weekly, daily, and hourly backups.
Execution Time	Set the execution time after the backup cycle is set to weekly or daily.
Start Time	Set the start time of the backup plan.

   For more information about backup policies, see *Backup Policy*.

   Retention Policy

Option	Description
Local Retention Policy	Set the local backup retention policy by quantity or by time.
Remote Retention Policy	Set this parameter after **Sync to Remote Backup Storage** is enabled. Set the remote backup retention policy to permanent retention, retention by quantity, or retention by time.

   For more information about retention policies, see *Backup Policy*.

7. Review the configuration and click **OK**.

What's next

After backups are generated, you can use backup data to restore the platform database. For more information, see *Data Recovery*.

11.5.3 View Backup Plans and Data

During backup plan execution, you can view the execution status of backup plans to monitor backup activity in real time. For more information, see *View Backup Plan Executions*.

During backup plan execution, backup data is generated. You can view backup data for protected resources. For more information, see *View Platform Database Backup Data*.

11.5.3.1 View Backup Plan Executions

View the Latest Backup Result

Procedure:

1. In the navigation pane, choose **Data Protection** > **Backup Plan**.
2. In the list, view the latest backup result from **Latest Backup Job Result**.

View Backup Job Details in a Backup Plan

You can view details of each backup job, including the backup object, backup mode, number of backup resources, total backup capacity, backup job result, backup duration, start and completion time, and backup details.

Procedure

1. In the navigation pane, choose **Data Protection** > **Backup Plan**.
2. On the **Backup Plan** page, click the platform database backup plan name to go to the details page.
3. On the backup plan details page, click **Backup Job**.
4. On the **Backup Job** tab, select the backup job that you want to view, and then click **View Backup Details**.
5. In the **Backup Job Details** dialog, view details of the backup job.

11.5.3.2 View Platform Database Backup Data

NexaSphere provides a unified entry for viewing backup data of protected resources so that important protected workloads are not missed.

Procedure

1. In the navigation pane, choose **Data Protection** > **Protected Resources**.
2. On the **Protected Resources** page, select the **Platform Database** tab.
3. View platform database backup data.

   The platform database backup data page is divided into tabs based on the storage location of the backup data:

   - Local Backup: Backup files are stored in local backup storage.
   - Remote Backup: Backup files are stored in remote backup storage.

After selecting the required location, you can view the name, management node IP address, version, capacity, backup storage path or synchronization status to remote or local backup storage, and creation time of the platform database backup.

11.6 Data Recovery

NexaSphere provides multiple data recovery methods for different recovery scenarios:

- You can restore an entire virtual machine from backup data to a specified location. For more information, see *Restore a VM from Backup*.
- You can create a new virtual machine directly from backup data. For more information, see *Create a VM from Backup*.
- You can restore the platform from backup data to a specified point in time. For more information, see *Restore the Platform Database from Backup*.

11.6.1 Restore a VM from Backup

If a virtual machine fails, you can restore the virtual machine from backup data to the latest backup point or an earlier backup point.

Prerequisites

- Make sure the virtual machine has at least one successfully generated backup data point.
- Before restoring a virtual machine in overwrite mode, shut down the virtual machine.
- To restore a virtual machine from remote backup data, synchronize the remote backup data to local backup storage first.

Procedure

1. In the navigation pane, choose **Data Protection** > **Protected Resources**.
2. On the **Protected Resources** page, select the **Virtual Machine** tab.
3. Select the backup data of the virtual machine to be restored, and then click **Restore**.
4. In the **Restore Virtual Machine** dialog, choose whether to automatically start the VM after restoration as needed.
5. Review the information and click **OK**.



Note:

Restoring a virtual machine overwrites the current VM data and restores VM data to the selected backup data point. Proceed with caution.

11.6.2 Create a VM from Backup

If a virtual machine fails, you can create a new virtual machine directly from backup data. This recovery method is non-destructive. It does not overwrite the existing VM, but creates a new VM from the backup data.

Prerequisites

- Make sure the virtual machine has at least one successfully generated backup data point.
- Make sure the platform has sufficient compute, storage, and network resources for the new VM .
- To create a VM from remote backup data, synchronize the remote backup data to local backup storage first.

Procedure

1. In the navigation pane, choose **Data Protection** > **Protected Resources**.
2. On the **Protected Resources** page, select the **Virtual Machine** tab.
3. Select the backup data of the virtual machine to be restored, and then click **Create VM**.
4. In the **Create VM from Backup** dialog, configure the following options:

   Backup Information

Option	Description
Backup Data	Displays the selected backup data.

   Basic Information

Option	Description
Name	Name of the virtual machine.
Quantity	Default is 1 and cannot be modified.
Group	Group to which the VM belongs. If this parameter is not set, the default group is used.
Running Location	Host or cluster where the VM runs.
Operating System	Displays the VM operating system recorded in the backup data.
High Availability	Automatic restart mechanism after abnormal VM shutdown. For more information, see *VM HA Policy*.

Option	Description
Power Status	Choose whether to automatically start the VM after it is created.

Hardware Information

Option	Description
CPU	Supports adjusting the total number of cores.
Memory	Supports adjusting the memory size.
Disk	Displays the disk configuration recorded in the backup data and cannot be modified.
NIC	Supports adjusting the port group, MAC address, IP address, DNS allocation mode, and security group. Click **Add NIC** to add a NIC to the VM. You can customize the address and features of the new NIC.

5. Review the configuration and click **OK**.

What's next

Some VM configurations depend on VMTools. After the VM is created, install VMTools to make these configurations take effect. For more information about VMTools, see *Install and Use VMTools*.

11.6.3 Restore the Platform Database from Backup

Restore the platform from backup data to a specified backup point.

Prerequisites

- Restoring the platform database requires restarting the management node, during which the management UI is unavailable. This process usually takes several minutes and does not affect your service resources.
- Resources deleted after the backup point become invalid data and cannot be restored to a normal state by restoring the platform database.

Procedure

1. In the navigation pane, choose **Data Protection** > **Protected Resources**.
2. On the **Protected Resources** page, select the **Platform Database** tab.

3. Select the storage location of the platform database backup file, such as **Local Backup** or **Remote Backup**.
4. In the platform database backup list, select the backup data to be restored, and then click **Actions** > **Restore**.
5. In the **Restore Platform Database** dialog, enter the root password of the platform database.
6. Review the information and click **OK**.

What's next

After the database is restored, all platform resources are restored to the state when the backup was created. Go to **Data Protection** > **Backup Storage**, select the backup storage, and then click **Actions** > **Scan Backup Data** to obtain current backup data information.

11.7 Backup Data Management

You can perform the following operations on backup data:

- *Synchronize Backup Data*
- *Change the Owner of VM Backup Data*
- *Export Platform Database Backup Data*
- *Delete Backup Data*

Synchronize Backup Data

You can synchronize one or more local backup data records to remote backup storage, or synchronize one or more remote backup data records to local backup storage to provide dual protection for backup data. This section uses VM backup data as an example to describe how to synchronize local backup data to remote backup storage.

Prerequisites

- Remote backup storage has been added to the platform in advance. For more information, see *Add Remote Backup Storage*.
- The selected backup data has not been synchronized before.
- The VM associated with the backup data and its disks are not encrypted.

Procedure

1. In the navigation pane, choose **Data Protection** > **Protected Resources**.
2. On the **Protected Resources** page, select the **Virtual Machine** tab.
3. Select **Local Backup**.

4. Select the VM whose backup data you want to synchronize, and then in the right-side **Backup Data** list, select the backup data to be synchronized.
5. Click **Bulk Actions** > **Sync to Remote Backup Storage**.
6. In the **Sync to Remote Backup Storage** dialog, specify one remote backup storage.
7. Review the information and click **OK**.

Change the Owner of VM Backup Data

You can change the owner of VM backup data.

Procedure

1. In the navigation pane, choose **Data Protection** > **Protected Resources**.
2. On the **Protected Resources** page, select the **Virtual Machine** tab.
3. Select the target VM, and then in the right-side **Backup Data** list, select the backup data whose owner you want to change.
4. Click **Bulk Actions** > **Change Owner**.
5. In the **Change Owner** dialog, specify the new owner.
6. Review the information and click **OK**.

Export Platform Database Backup Data

You can export platform database backup data to your local environment.

Procedure

1. In the navigation pane, choose **Data Protection** > **Protected Resources**.
2. On the **Protected Resources** page, select the **Platform Database** tab.
3. Select the storage location of the platform database backup file, such as **Local Backup** or **Remote Backup**.
4. In the platform database backup list, select the platform database backup to be exported, and then click **Actions** > **Export**.

Delete Backup Data

After you confirm that backup data of a protected resource is no longer needed, you can delete the backup data. This section uses VM backup data as an example.

Procedure

1. In the navigation pane, choose **Data Protection** > **Protected Resources**.
2. On the **Protected Resources** page, select the **Virtual Machine** tab.

3. Select the target VM, and then in the right-side **Backup Data** list, select the backup data to be deleted.
4. Click **Bulk Actions** > **Delete**.
5. In the **Delete Backup Data?** dialog, choose whether to also delete remote backup data as needed.
6. Review the information and click **OK**.

11.8 Backup Plan Management

You can perform the following operations on backup plans:

- *Enable/Disable Backup Plan*
- *Edit Name and Description*
- *Modify Basic Settings*
- *Modify Backup Settings*
- *Delete Backup Plan*

Enable/Disable Backup Plan

You can temporarily disable a periodically executed backup plan. A disabled backup plan is not deleted. Instead, it is paused for a period of time and does not run at the execution time specified by the backup cycle. Its associated protected resources also stop scheduled backup operations until the backup plan is enabled again.

Procedure

1. In the navigation pane, choose **Data Protection** > **Backup Plan**.
2. On the **Backup Plan** page, select the backup plan to be disabled or enabled, and then click **Actions** > **Disable/Enable**.

Edit Name and Description

You can modify the name or description of a backup plan at any time.

Procedure

1. In the navigation pane, choose **Data Protection** > **Backup Plan**.
2. On the **Backup Plan** page, select the target backup plan, and then click **Actions** > **Edit Name and Description**.
3. In the **Edit Name and Description** dialog, enter the new name or description, and then click **OK**.

Modify Basic Settings

You can modify the basic settings of a backup plan at any time, or add virtual machines to the backup plan.

Procedure

1. In the navigation pane, choose **Data Protection** > **Backup Plan**.
2. On the **Backup Plan** page, select the target backup plan, and then click **Actions** > **Modify Basic Settings**.
3. In the **Modify Basic Settings** dialog, modify basic information, backup storage information, or backup object information as needed.
4. Review the configuration and click **OK**.

Modify Backup Settings

You can modify backup settings of a backup plan at any time, such as the backup policy, retention policy, and advanced settings.

For more information about VM backup policies, see *Backup Policy*. For more information about platform database backup policies, see *Backup Policy*.

Procedure

1. In the navigation pane, choose **Data Protection** > **Backup Plan**.
2. On the **Backup Plan** page, select the target backup plan, and then click **Actions** > **Modify Backup Settings**.
3. In the **Modify Backup Settings** dialog, modify the backup policy, retention policy, or advanced settings as needed.
4. Review the configuration and click **OK**.

Delete Backup Plan

If a backup plan is no longer needed, you can delete it. After the backup plan is deleted, associated resources no longer perform scheduled backups. To continue protecting specific resources, create a new backup plan for them.

Procedure

1. In the navigation pane, choose **Data Protection** > **Backup Plan**.
2. On the **Backup Plan** page, select the target backup plan, and then click **Actions** > **Delete**.

11.9 Backup Storage Management

You can perform the following operations on backup storage:

- *Enable/Disable/Reconnect Backup Storage*
- *Scan Backup Data*
- *Clean Up Backup Storage Data*
- *Edit Name and Description*
- *Modify Backup Storage Configuration*
- *Update Backup Storage Password*
- *Delete Backup Storage*

Enable/Disable/Reconnect Backup Storage

You can manage backup storage status, including enabling, disabling, and reconnecting backup storage.

Disabling backup storage might affect backup plan execution. Note the following:

- If a backup plan specifies only one backup storage, the backup plan fails after the backup storage is disabled.
- If a backup plan specifies two backup storage resources, after one backup storage resource , the primary one, is disabled, the backup plan automatically switches to the other backup storage resource, the standby one, to perform backups.

Procedure

1. In the navigation pane, choose **Data Protection** > **Backup Storage**.
2. On the **Backup Storage** page, select the target backup storage, and then click **Actions** > **Enable/Disable/Reconnect**.

Scan Backup Data

You can scan backup data that has not been deleted from backup storage. Deleted backup data cannot be restored by scanning. If you manually add backup data to backup storage and want to use the backup data, you can also scan the backup data on the backup storage.

Procedure

1. In the navigation pane, choose **Data Protection** > **Backup Storage**.
2. On the **Backup Storage** page, select the target backup storage, and then click **Actions** > **Scan Backup Data**.

Clean Up Backup Storage Data

You can clean up invalid backup data that has been permanently deleted from backup storage to release storage space.

Procedure

1. In the navigation pane, choose **Data Protection** > **Backup Storage**.
2. On the **Backup Storage** page, select the target backup storage, and then click **Actions** > **Cleanup Data**.

Edit Name and Description

You can modify the name or description of backup storage at any time.

Procedure

1. In the navigation pane, choose **Data Protection** > **Backup Storage**.
2. On the **Backup Storage** page, select the target backup storage, and then click **Actions** > **Edit Name and Description**.
3. In the **Edit Name and Description** dialog, enter the new name or description, and then click **OK**.

Modify Backup Storage Configuration

You can modify configuration information for added backup storage.

Procedure

1. In the navigation pane, choose **Data Protection** > **Backup Storage**.
2. On the **Backup Storage** page, select the target backup storage, and then click **Actions** > **Modify Configuration**.
3. In the **Modify Configuration** dialog, modify basic information or configuration information as needed.
4. Review the information and click **OK**.

Update Backup Storage Password

You can modify the password of added backup storage.

Procedure

1. In the navigation pane, choose **Data Protection** > **Backup Storage**.

2. On the **Backup Storage** page, select the target backup storage, and then click **Actions** > **Update Password**.
3. In the **Update Password** dialog, enter the new password.
4. Review the information and click **OK**.

Delete Backup Storage

Deleting backup storage is a high-risk operation. This operation stops all backup plans on the backup storage, and related backup data can no longer be viewed.

Procedure

1. In the navigation pane, choose **Data Protection** > **Backup Storage**.
2. On the **Backup Storage** page, select the target backup storage, and then click **Actions** > **Delete**.

12 Bare Metal Management

12.1 Overview

NexaSphere Bare Metal Management provides dedicated physical servers for critical applications that require high performance and stability.

After racking the servers and completing the basic preparation, administrators can add bare metal chassis in batches and create bare metal instances from the UI. Bare metal templates support unattended batch operating system installation, while business networks can be configured for bare metal instances to meet different application requirements.

Core Concepts

- Bare Metal Template: With bare metal templates, preconfigured files can be quickly generated to achieve unattended bulk OS installation for bare metal instances.
- Bare Metal Cluster: A bare metal cluster consists of bare metal chassis.
- Deployment Server: An independent server used for providing PXE services and console proxies for bare metal chassis.
- Bare Metal Chassis: Bare metal chassis is used to create bare metal instances and can be universally identified by the BMC interface and IPMI configuration.
- Bare Metal Instance: A virtualized instance of bare metal chassis.

Advantages

Bare Metal Management provides the following advantages:

- **Dedicated physical resources**

  Provides dedicated physical servers for critical applications to ensure high performance and stability.
- **Efficient batch deployment**
  - Adds bare metal chassis in batches from the UI. Chassis can be added manually or imported from a template, allowing you to configure IPMI addresses for multiple chassis at once and quickly build bare metal clusters.
  - Uses bare metal templates for unattended operating system installation to improve operational efficiency.
- **Flexible operating system support**

Supports mainstream Linux distributions, including the RHEL/CentOS, Debian/Ubuntu, and SUSE/openSUSE families, as well as customized platform operating systems and custom installation requirements.

- **Highly available deployment architecture**

  Supports highly available environments with two management nodes. A dedicated deployment server is recommended for each bare metal cluster to avoid single points of failure and DHCP conflicts.

Typical Use Cases

Bare Metal Management is suitable for the following use cases:

- **Highly regulated environments**

  Industries such as finance and securities have stringent requirements for deployment compliance and data security. Dedicated resources and data isolation help make operations auditable and traceable.

- **High-performance computing**

  Supercomputing, genome sequencing, and similar workloads require high compute performance, stability, and responsiveness. Deploying bare metal clusters avoids virtualization overhead and meets these performance requirements.

- **Mission-critical databases**

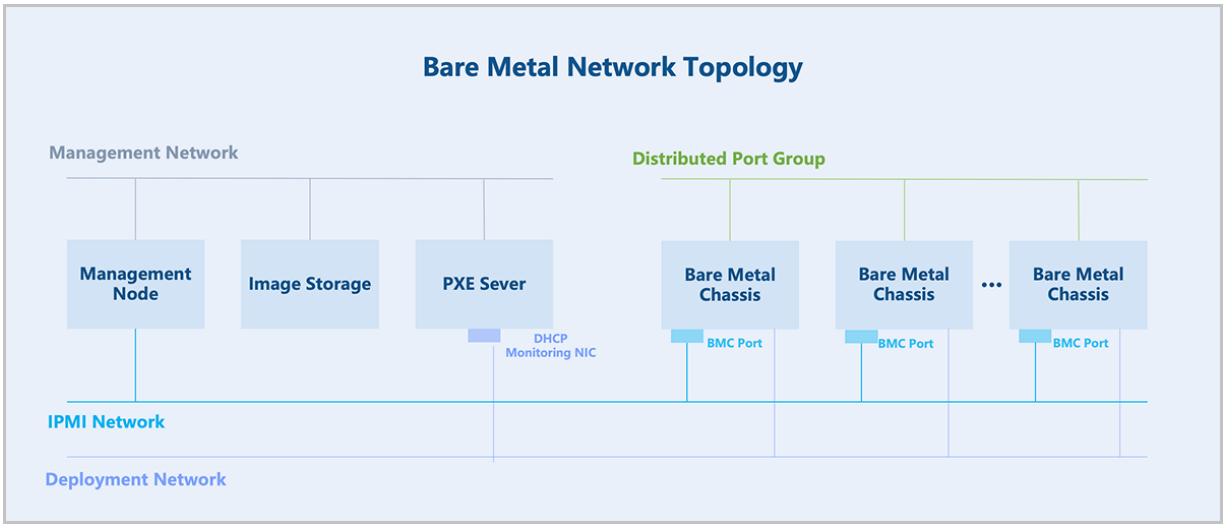
  Databases that cannot run in virtual machines can use dedicated physical servers with isolated networks and predictable performance.

12.2 Preparation

Network Planning

Plan the physical network environment, including the deployment network, IPMI network, management network, and business network, as shown in *Figure 12-1: Bare Metal Network Topology*.

Figure 12-1: Bare Metal Network Topology



1. Plan the deployment network.

   Make sure that the PXE NIC on each bare metal chassis can communicate with the DHCP listening NIC on the deployment server (PXE server) through the deployment network.

   - Access the BIOS of each bare metal chassis and enable PXE on the NIC connected to the deployment network.

     For some server models, make sure that the PXE NIC is the first boot NIC, or disable PXE on all NICs with a higher boot priority.
   - You do not need to manually change the BIOS boot order, which typically prioritizes booting from the hard disk. The system automatically triggers a one-time PXE boot.
   - To avoid a single point of failure, use a dedicated deployment server.
   - The DHCP listening NIC on the deployment server must be a dedicated NIC with an IP address so that it can provide stable DHCP services.
   - Make sure that no DHCP service other than the one on the deployment server is available on the deployment network.
   - Plan the IP address assigned to each bare metal instance according to your production environment.

2. Configure IPMI for each bare metal chassis and plan the IPMI network.

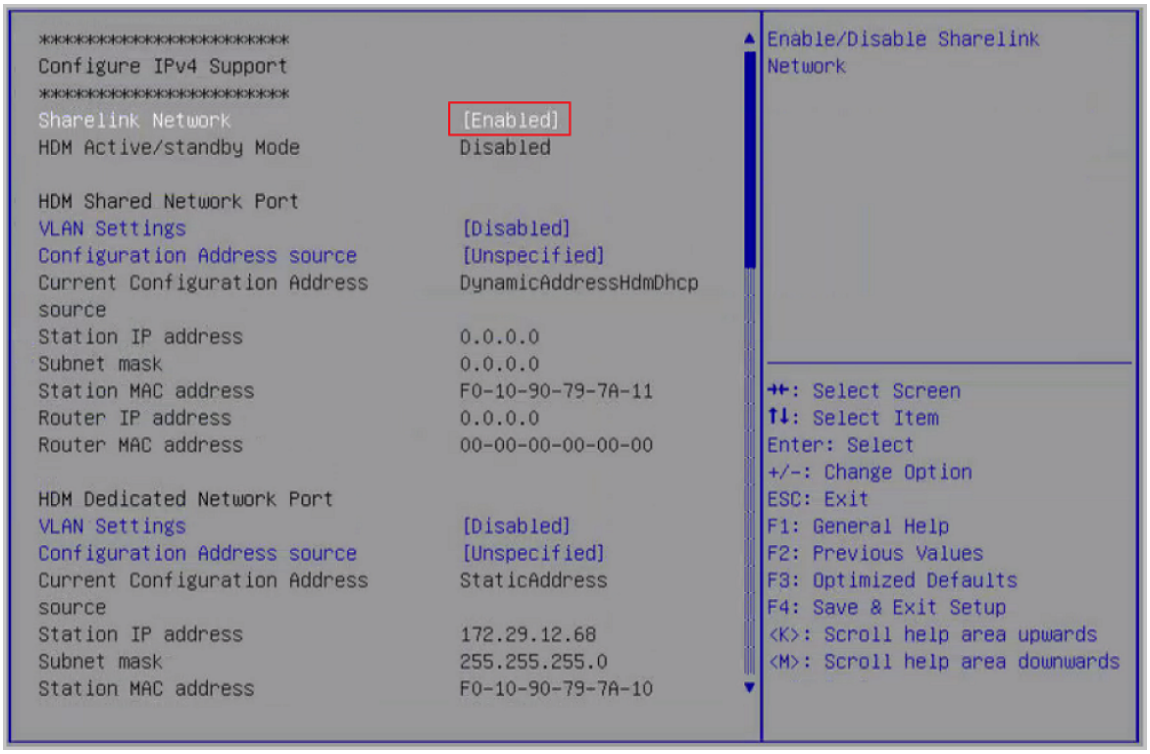
   Consider the following requirements:

   - For out-of-band management, each bare metal chassis must have a BMC interface. Configure the IPMI address, port, username, and password for each chassis. Make sure

that the management nodes can communicate with the BMC interfaces through the IPMI network.

- If a bare metal chassis supports the BMC ShareLink Management Port, disable it in the BIOS so that the chassis has only one IPMI address.

Figure 12-2: Disable the BMC ShareLink Management Port



3. Plan the management network and business networks.
   - Make sure that the image storage and deployment servers can communicate with the management nodes through the management network.
   - To configure business networks for bare metal instances, attach the required distributed switches to the bare metal cluster in advance.

Software Preparation

Prepare ISO images in the image storage for installing operating systems on bare metal instances.

- Bare metal instances support customized platform operating systems and mainstream Linux distributions, including the RHEL/CentOS, Debian/Ubuntu, and SUSE/openSUSE families.
- Only ISO images are supported.

12.3 Quick Start Guide

Follow this workflow to create bare metal instances:

1. (Optional) Add a custom bare metal template.

   Bare metal templates generate preconfigured files for unattended batch operating system installation. You can use a system preconfigured template or upload a custom template. For more information, see *Add a Custom Template*.

2. Create a bare metal cluster.

   A bare metal cluster provides a dedicated management boundary for bare metal chassis. For more information, see *Create a Bare Metal Cluster*.

3. Add bare metal chassis.

   Add chassis manually or import them from a template. For more information, see *Manually Add Bare Metal Chassis* and *Import Bare Metal Chassis from a Template*.

4. Create a bare metal instance.

   For more information, see *Create a Bare Metal Instance*.

12.4 Bare Metal Templates

12.4.1 Template Syntax Rules

Bare metal templates include system variables and custom variables to support unattended deployment scenarios of varying complexity.

System preconfigured templates contain only system variables. Custom templates contain the basic system variables and can also include variables that you define.

System variable example (use uppercase letters and underscores as separators):

```
REPO_URL
# Installation source URL created from the selected ISO image
# Can be commented out and manually specified by using --url

USERNAME
# System username
# The default username is root for RHEL/CentOS and SUSE/openSUSE
 systems; only a password is required
# A custom username and password are required for Debian/Ubuntu
 systems

PASSWORD
# Password for the system user

NETWORK_CFGS
# NIC UUID and network UUID entered in the UI
```

```
# Replaced after automatic or manual IP address assignment

FORCE_INSTALL
# Whether to overwrite existing disk data automatically during
 deployment

PRE_SCRIPTS
# Scripts to run before deployment

POST_SCRIPTS
# Scripts to run after deployment
```

Custom variable example (use lowercase letters and underscores as separators):

```
hostname
# Hostname

keyboard
# Keyboard

timezone
# Time zone
```

Use the syntax rules for the corresponding template type:

- For kickstart templates, see the Red Hat documentation.
- For preseed or autoinstall templates, see the Ubuntu documentation.
- For autoyast templates, see the SUSE documentation.

12.4.2 System Preconfigured Templates

NexaSphere provides the following system preconfigured templates:

Template Name	Description
cloud_host_x86_64_v3	Used for unattended deployment of NexaSphere management nodes through Bare Metal Management.
cloud_expert_x86_64_v2	Used for unattended deployment of NexaSphere in streamlined expert mode through Bare Metal Management.
centos_7_x86_64_mini_v1	Used for unattended deployment of CentOS 7 through Bare Metal Management.
centos_7_aarch64_min_v1	
kylin_10_x86_64_min_v1	Used for unattended deployment of Kylin V10 through Bare Metal Management.
kylin_10_aarch64_min_v1	
openEuler_20_aarch64 _min_v1	Used for unattended deployment of openEuler 20 through Bare Metal Management.
opensuse_15_x86_64_m ini_v1	Used for unattended deployment of openSUSE 15 through Bare Metal Management.

Template Name	Description
ubuntu_16_x86_64_mini_v2	Used for unattended deployment of Ubuntu 16 through Bare Metal Management.
ubuntu_18_x86_64_mini_v1	Used for unattended deployment of Ubuntu 18 through Bare Metal Management.
ubuntu_20_live_server_x86_64_mini_v1	Used for unattended deployment of Ubuntu 20 and Ubuntu 22 through Bare Metal Management.

12.4.3 Add a Custom Template

Upload a UTF-8 encoded custom bare metal template file for complex unattended deployment scenarios.

Prerequisites

- The custom template file must not exceed 50 KB.
- The custom template file must follow the syntax rules for its template type. For more information, see *Template Syntax Rules*.

Procedure

1. In the navigation pane, choose **Inventory** > **Bare Metal Management**.
2. Select the root node.
3. On the root node details page, click **Customized Configuration**.
4. On the **Customized Configuration** tab, click **Add Bare Metal Template**.
5. On the **Add Bare Metal Template** dialog, configure the following options:

Option	Description
Name	Enter a name for the bare metal template.
Description	Enter a brief description for the bare metal template.
Operating System	Select the operating system for unattended deployment. Customized platform operating systems and mainstream Linux distributions are supported, including the RHEL/CentOS, Debian/Ubuntu, and SUSE/openSUSE families.

Option	Description
Template Type	Select a template type that matches the selected operating system. • Customized platform operating systems: Select kickstart. • RHEL/CentOS family: Select kickstart. • Debian/Ubuntu family: Select preseed. • Ubuntu Live: Select autoinstall. • SUSE/openSUSE family: Select autoyast.
Template Import	Upload a UTF-8 encoded custom template file. **Note:** Write the custom bare metal template according to the syntax rules for the selected template type.

6. Review the configuration and click **OK**.

12.5 Bare Metal Clusters

12.5.1 Create a Bare Metal Cluster

Prerequisites

- Attach a deployment server to the bare metal cluster to provide PXE services for bare metal instances in the cluster.
- Each bare metal cluster supports only one deployment server. A deployment server can be attached to multiple bare metal clusters.
- You can attach distributed switches to a bare metal cluster to provide network services for bare metal instances in the cluster.

Procedure

1. In the navigation pane, choose **Inventory** > **Bare Metal Management**.
2. Select the target data center.
3. On the data center details page, click **Bare Metal Management**.

4. On the **Bare Metal Management** tab, choose **Bare Metal Cluster** > **New Bare Metal Cluster**.
5. In the **New Bare Metal Cluster** dialog, configure the following options:

Option	Description
Name	Enter a name for the bare metal cluster.
Description	Enter a brief description for the bare metal cluster.
Attach Deployment Server	Select whether to attach a deployment server.
Server Source	Select whether to create a deployment server or use an existing one. If you select **New Deployment Server**, configure the deployment server options. If you select **Existing Deployment Server**, select a deployment server.
Deployment Server Name	If you select **New Deployment Server**, enter a name for the deployment server.
DHCP Listening NIC	If you select **New Deployment Server**, enter the device name of the NIC on the deployment server that is connected to the deployment network. **Note:** • The NIC must be connected to the deployment network of the bare metal chassis and have an IP address. • No other DHCP service can be available on the network connected to this NIC.
Storage Path	If you select **New Deployment Server**, enter a local directory on the deployment server for storing PXE deployment images.

Option	Description
Deployment Server IP	If you select **New Deployment Server**, enter the IP address of the deployment server.
SSH Port	If you select **New Deployment Server**, enter the SSH port. The default value is 22.
Username	If you select **New Deployment Server**, enter a username. The default username is root. You can also enter a regular user.
Password	If you select **New Deployment Server**, enter the password for the specified user.
DHCP Start IP	If you select **New Deployment Server**, enter the start IP address of the range scanned by the DHCP service.
DHCP End IP	If you select **New Deployment Server**, enter the end IP address of the range scanned by the DHCP service. **Note:** If you leave this field empty, the system detects the network based on the NIC IP address and excludes IP addresses that are already in use.

6. Review the configuration and click **OK**.

12.5.2 Manage Deployment Servers

Procedure

1. In the navigation pane, choose **Inventory** > **Bare Metal Management**.
2. Select the target bare metal cluster.
3. On the **Overview** tab of the bare metal cluster, view the configuration of the attached deployment server.
4. On the bare metal cluster details page, click **Actions**.

- To attach a deployment server, select **Attach Deployment Server**.
- To reconnect a deployment server, select **Reconnect Deployment Server**.
- To detach the deployment server from the bare metal cluster, select **Detach Deployment Server**.

12.5.3 Delete a Bare Metal Cluster

Procedure

1. In the navigation pane, choose **Inventory** > **Bare Metal Management**.
2. Select the target bare metal cluster.
3. On the bare metal cluster details page, choose **Actions** > **Delete**.

Result

Note:
Deleting a bare metal cluster also deletes all bare metal chassis and bare metal instances in the cluster. Proceed with caution.

12.6 Bare Metal Chassis

12.6.1 Add Bare Metal Chassis

12.6.1.1 Manually Add Bare Metal Chassis

Add one or more bare metal chassis by specifying an IPMI address or an IPMI address range.

Prerequisites

Make sure that the bare metal chassis are connected according to the planned network topology. For more information, see *Preparation*.

Procedure

1. In the navigation pane, choose **Inventory** > **Bare Metal Management**.
2. Right-click the target bare metal cluster and select **Add Bare Metal Chassis**.
3. In the **Select Bare Metal Chassis Addition Type** dialog, select **Manual Addition**.
4. Click **Next**.
5. In the **Add Bare Metal Chassis** dialog, configure the following options:

Option	Description
Name	Enter a name for the bare metal chassis.
Description	Enter a brief description for the bare metal chassis.
Bare Metal Cluster	Select the bare metal cluster to which the chassis belongs.
Addition Method	Select how to add IPMI addresses. You can specify an IPMI address or an IPMI address range.
IPMI Address/Range	Enter an IPMI address or the start and end addresses of an IPMI address range.
Port	Enter the IPMI port. The default port is 623.
Username	Enter the IPMI username.
Password	Enter the IPMI password.
Reboot Bare Metal Chassis	Select whether to reboot the bare metal chassis to automatically obtain hardware information. This option is cleared by default.

6. Review the configuration and click **OK**.

12.6.1.2 Import Bare Metal Chassis from a Template

Enter bare metal chassis information in a CSV configuration file and upload the file to import the chassis.

Prerequisites

- Make sure that the bare metal chassis are connected according to the planned network topology. For more information, see *Preparation*.
- The template includes a header row and an example row. Delete or overwrite the example row when editing the template.
- Options marked with * are required.
- You can enter a single IPMI address or an IPMI address range.

For IP address ranges, use commas to separate ranges and use ^ to exclude a range. Example: *127.0.0.1-127.0.0.10,^127.0.0.2-127.0.0.3*.

- For the option to reboot bare metal chassis:
  - To automatically reboot the chassis and obtain hardware information, enter YES, Yes, yes, Y, or y.
  - To reboot the chassis manually, enter NO, No, no, N, or n, or leave the field empty.

Procedure

1. In the navigation pane, choose **Inventory** > **Bare Metal Management**.
2. Right-click the target bare metal cluster and select **Add Bare Metal Chassis**.
3. In the **Select Bare Metal Chassis Addition Type** dialog, select **Template Import**.
4. Click **Next**.
5. In the **Add Bare Metal Chassis** dialog, configure the following options:

Option	Description
Template	Click **Download Template** and enter the bare metal chassis information in the required format.
Upload Template File	Upload the completed template file.

6. Click **OK**.

12.6.2 Bare Metal Chassis Status Management

The deployment status, status, and power status of bare metal chassis are described in the following tables:

Deployment Status	Description
Unknown Hardware Info	The system has not obtained hardware information from the bare metal chassis. **Note:** • Possible cause: The option to reboot the bare metal chassis was not selected when the chassis was added. • Solution: Reboot the bare metal chassis manually or manually obtain its hardware information.

Deployment Status	Description
PXE Booting	The deployment server (PXE server) is instructing the bare metal chassis to boot from the PXE NIC and is assigning a dynamic IP address.
PXE Boot Failed	The bare metal chassis failed to boot from the PXE NIC. **Note:** Check the following requirements: • Make sure that no DHCP service other than the one on the deployment server is available on the deployment network. • Make sure that PXE is enabled in the BIOS for the NIC connected to the deployment network. For some server models, make sure that the PXE NIC is the first boot NIC, or disable PXE on all NICs with a higher boot priority. • Make sure that the boot mode of the bare metal chassis is Legacy.
Available	The bare metal chassis can be used to create a bare metal instance.
Allocated	The bare metal chassis has been used to create a bare metal instance.
Rebooting	The bare metal chassis is restarting.

Status	Power Status
Status can be: • Enabled • Disabled	Power status can be: • Power On • Rebooting • Power Off

12.6.3 Access a Bare Metal Chassis

Procedure

1. In the navigation pane, choose **Inventory** > **Bare Metal Management**.
2. Select the target bare metal chassis.
3. On the bare metal chassis details page, click **Launch Console**.
   The system redirects you to the IPMI management interface of the bare metal chassis.
4. On the IPMI management interface, enter the configured IPMI username and password to log in.

12.6.4 Delete a Bare Metal Chassis

Procedure

1. In the navigation pane, choose **Inventory** > **Bare Metal Management**.
2. Select the target bare metal chassis.
3. On the bare metal chassis details page, choose **Actions** > **Delete**.

Result

Note:

Deleting a bare metal chassis also deletes the bare metal instance created from that chassis.

12.7 Bare Metal Instances

12.7.1 Create a Bare Metal Instance

Procedure

1. In the navigation pane, choose **Inventory** > **Bare Metal Management**.
2. Select the target bare metal chassis.
3. On the bare metal chassis details page, choose **Actions** > **New Bare Metal Instance**.
4. In the **New Bare Metal Instance** dialog, configure the following basic information:

Option	Description
Name	Enter a name for the bare metal instance.
Description	Enter a brief description for the bare metal instance.
Bare Metal Chassis	Displays the bare metal chassis on which the instance will be created.
Platform	Only Linux is supported.
Image	Select an image to install the operating system. **Note:** • Only ISO images are supported.

Option	Description
	• Upload the image to a standalone image repository in advance. • Customized platform operating systems and mainstream Linux distributions are supported, including the RHEL/CentOS, Debian/Ubuntu, and SUSE/openSUSE families.
Bare Metal Template	Select a bare metal template to generate a preconfigured file for unattended batch operating system installation. For more information about template syntax and system preconfigured templates, see *Template Syntax Rules* and *System Preconfigured Templates*.
Auto Overwrite Data	Select whether to automatically overwrite data. **Note:** • If enabled, the system automatically overwrites existing data on the disks during deployment. • If disabled, the operating system installation might pause. You must then access the console and configure the disks manually.

5. Click **Next**.
6. Configure the following system and network information:

   You can configure system and network information for bare metal instances individually or in batches.

Option	Description
Operating System	Displays the selected operating system.
Username	Enter a username. For batch configuration, the default username is root.
Password	Enter a password.
Network Device	Configure a business network for the bare metal instance.
Device Type	Select NIC or NIC Bond.
NIC Bond Name	If you select NIC Bond as the device type, enter a name for the NIC bond.
NIC Bond Type	If you select NIC Bond as the device type, select Mode 1 (Active-Backup) or Mode 4 (LACP).
NIC	Select a NIC.
Distributed Port Group	Select a distributed port group.

7. Review the configuration and click **OK**.

12.7.2 Bare Metal Instance Status Management

The state and deployment status of bare metal instances are described in the following table:

State	Deployment Status
State can be: • Created • Powering On • Running • Stopped • Rebooting • Deleted • Error • Unknown	Deployment status can be: • Unprovisioned • Provisioning • Deployed

12.7.3 Delete a Bare Metal Instance

Prerequisites

(Optional) The platform provides deletion protection for bare metal instances. You can control how resources are deleted by configuring the deletion policy in system parameters. By default, the bare metal instance deletion policy is Never Delete. Deleted resources are moved to the recycle bin and retained until you manually expunge them. For more information, see *Platform Policy*.

Procedure

1. In the navigation pane, choose **Inventory** > **Bare Metal Management**.
2. Select the target bare metal instance.
3. On the bare metal instance details page, choose **Actions** > **Move to Recycle Bin**.

   To delete multiple bare metal instances, go to the **Bare Metal Management** tab on the data center details page. Select the bare metal instances, and then choose **Bulk Action** > **Move to Recycle Bin**.

   The button label depends on the **Bare Metal Instance Deletion Policy**. When the policy is Immediate Deletion, the button is labeled "Delete". When the policy is Never Delete, the button is labeled "Move to Recycle Bin".
4. In the confirmation dialog, acknowledge the risk and click **OK**.

Result

Note:
Deleting a bare metal instance also powers off the corresponding bare metal chassis and might interrupt services. Proceed with caution.

13 Platform Configuration

13.1 Overview

NexaSphere Platform Configuration provides platform-wide settings that help administrators maintain platform time, customize the user interface, control default platform behavior, and configure external log delivery and email services.

You can use Platform Configuration to perform the following operations:

- Configure time sources and synchronize platform time.
- Send management node logs or platform operation logs to external logging systems.
- Configure email servers for sending alarm messages.
- Customize or restore the default theme and appearance.
- View, modify, and restore system parameters.

13.2 Time Management

Manage platform time and configure time sources. After configuration, all hosts on the platform synchronize their time with the specified time sources.

Related Definitions

Time Management uses the following concepts:

- Internal time source: Uses a management node or host as the time source for the platform. Other nodes on the platform synchronize their time with the internal time source.
- External time source: Uses a node outside the platform as the time source. Platform nodes synchronize their time with it either directly or through an internal time source.
- Time synchronization: The process of adjusting the time of a platform node to match the time source.

How It Works

The platform uses Network Time Protocol and clock adjustment algorithms to keep local clocks synchronized with time sources. The basic process is as follows:

1. Select a time source: Select a configured internal or external time source as the time reference.
2. Obtain timestamps: Each platform node communicates with the time source to obtain timestamp information.

3. Calculate the time difference: Each platform node calculates the difference between its local time and the time source.
4. Adjust the local clock: Each platform node adjusts its local clock based on the calculated time difference.
5. Synchronize periodically: Each platform node periodically communicates with the time source to prevent clock drift and maintain synchronization.

Advantages

Time Management provides the following advantages:

- Precision: Uses accurate clock adjustment algorithms for precise time synchronization.
- Automation: Periodically calibrates time without manual intervention.
- Reliability: Supports multiple time sources to improve synchronization reliability and stability.
- Visibility: Displays the relationship between time sources and platform time configuration in a topology, together with the current platform time and time zone.

Use Cases

Time Management is commonly used in the following scenarios:

- Network management: Provides accurate timestamps for logs collected from different network devices to support troubleshooting.
- Billing systems: Provides consistent timekeeping to ensure accurate billing records.
- Coordinated processing: Provides a common time reference so that multiple systems process the same complex event in the correct order.

13.2.1 Modify Time Source Configuration

Context

The following time synchronization modes are supported:

- Internal: Configure only internal time sources. Other platform nodes synchronize with the internal time sources. You can add up to two internal time sources.
- External: Configure only external time sources. Platform nodes synchronize directly with the external time sources. You can add up to two external time sources.
- Internal and External: Configure both internal and external time sources. The internal time sources first synchronize with the external time sources, and then other platform nodes

synchronize with the internal time sources. You can add up to two internal time sources and two external time sources.

Procedure

1. In the navigation pane, choose **System Management** > **Time Configuration**.
2. On the **Time Configuration** page, click **Modify Configuration**.
3. In the **Modify Configuration** dialog, configure the following options:

Option	Description
NTP Mode	Select Internal, External, or Internal and External.
Time Source	Add time sources for the selected NTP mode. You can add up to two internal time sources and two external time sources. An external time source can be specified by an IP address or domain name.



Note:

- Before modifying time sources, make sure that the management nodes are connected.
- Modifying time sources might cause inconsistent platform node times, deviations or errors in monitoring data, and interruptions to running tasks. Evaluate the impact before making changes.

4. Review the configuration and click **OK**.

13.2.2 Synchronize Time

If the system time differs significantly from the time source, force a synchronization to correct the time immediately instead of waiting for gradual adjustment.

Procedure

1. In the navigation pane, choose **System Management** > **Time Configuration**.
2. On the **Time Configuration** page, click **Sync Time**.

13.3 Log Server Management

Log server management sends management node logs or platform operation logs to external logging systems for centralized retention, search, and analysis.

Select an output based on the type of logs that you want to send:

Log Type	Purpose	Supported Outputs
Management Node Log	Centrally retains management node logs for troubleshooting.	Syslog.
Platform Operation Log	Sends platform operation logs to an external logging system for centralized search and analysis.	Syslog, Elasticsearch, Forward, Kafka, and Loki.

13.3.1 Add a Log Server

Add a log server to send management node logs or platform operation logs to an external logging system.

Prerequisites

- The current user must have admin privileges.
- An external logging system has been deployed to receive logs.
- The management node can communicate with the external logging system over the network.

Procedure

1. In the navigation pane, choose **System Management** > **Log Server**.
2. On the **Log Server** page, click **Add Log Server**.
3. Configure the log server.

Option	Description
Name	Set the log server name.
Description	Enter a brief description for the log server.
Log Type	Select the type of logs to send: • **Management Node Log**: Sends management node logs. • **Platform Operation Log**: Sends platform operation logs.

Option	Description
Output	Select a log output: • Management node logs can be sent to Syslog. • Platform operation logs can be sent to Syslog, Elasticsearch, Forward, Kafka, or Loki.

4. Configure the output.

   For details about the options for each output, see *Log Server Configuration Options*.

5. Optional: Click **Test Connection** to test the network connection between the platform and the log server.

6. Review the configuration and click **OK**.

Result

The log server is added and displayed in the log server list.

13.3.2 Log Server Configuration Options

The available log server options depend on the log type and output.

Syslog

Option	Description
Address	Enter the IP address, domain name, or URL of the Syslog server.
Port	Enter the Syslog service port. The default port is 514.
Protocol	Select UDP or TCP.
Log Facility	This option is available for management node logs. Select a log facility from LOCAL0 through LOCAL7. The value must match the configuration on the Syslog server.
Log Severity Level	This option is available for management node logs. Select ALL, TRACE, DEBUG, INFO, WARN, ERROR, or FATAL.

Elasticsearch

Option	Description
Address	Enter the IP address, domain name, or URL of the Elasticsearch service.
Port	Enter the Elasticsearch service port. The default port is 9200.
Index	Enter the index used to store platform operation logs.
TLS	Specify whether to encrypt the connection with TLS.
Username	Enter the username for the Elasticsearch service.
Password	Enter the password for the Elasticsearch service.

Forward

Option	Description
Address	Enter the IP address, domain name, or URL of the Forward service.
Port	Enter the Forward service port. The default port is 24224.

Kafka

Option	Description
Address	Enter the IP address, domain name, or URL of the Kafka service.
Port	Enter the Kafka service port. The default port is 9092.
Topics	Enter the Kafka topic used to receive platform operation logs.

Loki

Option	Description
Address	Enter the IP address, domain name, or URL of the Loki service.
Port	Enter the Loki service port. The default port is 3100.
Labels Job	Set the log label written to Loki. The default value is `job=operator`.

Option	Description
TLS	Specify whether to encrypt the connection with TLS.
Username	Enter the username for the Loki service.
Password	Enter the password for the Loki service.

13.3.3 Edit, Test, and Delete Log Servers

From the log server list, you can edit log server information, test network connections, and delete log servers that are no longer needed.

In the navigation pane, choose **System Management** > **Log Server**. Use the action menu in a resource row to manage an individual log server. To perform a batch operation, select one or more log servers and click **Actions**.

Operation	Description
Edit	Change the name and description of a log server.
Test Connection	Test the network connection between the platform and one or more log servers.
Delete	Delete one or more log servers that are no longer needed.

13.4 Email Server Management

Email servers provide SMTP email delivery for sending alarm messages to email endpoints.

After adding an email server, you can create an email endpoint in the alarm service and select the server to send alarm messages.

13.4.1 Add an Email Server

Procedure

1. In the navigation pane, choose **System Management** > **Email Server**.
2. On the **Email Server** page, click **Add Email Server**.
3. Configure the email server.

Option	Description
Name	Set the email server name.

Option	Description
Description	Enter a brief description for the email server.
Username	Enter the email server username.
Password	Enter the password for the username.
Email Server Type	The default protocol is SMTP.
SMTP Server	Enter the email server address.
Encryption Type	Select whether to use an encrypted connection for the email server port. STARTTLS, SSL/TLS, and No Encryption are supported.
SMTP Port	Set the email server port.

4. Review the configuration and click **OK**.

13.4.2 Edit, Enable or Disable, and Delete Email Servers

From the email server list, you can edit email server information, enable or disable email servers, and delete email servers that are no longer needed.

In the navigation pane, choose **System Management** > **Email Server**. On the **Email Server** page, use the action menu in a resource row to manage the email server.

Operation	Description
Edit Name and Description	Change the name and description of the email server.
Enable/Disable	Enable or disable the email server.
Delete	Delete an email server that is no longer needed.

13.5 Theme and Appearance

13.5.1 Customize Theme and Appearance

Procedure

1. In the navigation pane, choose **System Management** > **Theme and Appearance**.
2. On the **Theme and Appearance** page, click the edit button to customize the theme and appearance.

3. For **Theme**, select one of the eight available interface theme colors.
4. For **Browser Title**, configure the following options.

Option	Description
Favicon	Upload a browser icon in `.ico` format. The file size cannot exceed 2 MB.
Chinese Title	Enter a Chinese browser title of up to 25 characters.
English Title	Enter an English browser title of up to 25 characters.

5. For **Login Interface Title**, configure the following options.

Option	Description
Logo	Upload a login interface logo in `.jpg`, `.jpeg`, `.png`, or `.svg` format. The image must not exceed 250 x 70 px or 2 MB.
Chinese Title	Enter a Chinese login interface title of up to 25 characters.
English Title	Enter an English login interface title of up to 25 characters.

6. For **Platform Interface Title**, configure the following options.

Option	Description
Logo	Upload a platform interface logo in `.jpg`, `.jpeg`, `.png`, or `.svg` format. The image must not exceed 110 x 40 px or 2 MB. **Note:** On a dark background, use a white or light-colored logo.
Chinese Title	Enter a Chinese platform interface title of up to 25 characters.

Option	Description
English Title	Enter an English platform interface title of up to 25 characters.
Font Size	Select Large, Medium, or Small for the platform interface title. The default value is Medium.

13.5.2 Restore Default Theme and Appearance

Clear the current customizations and restore the default theme and appearance.

Procedure

1. In the navigation pane, choose **System Management** > **Theme and Appearance**.
2. On the **Theme and Appearance** page, click **Reset to Default Settings**.

13.6 System Parameters

NexaSphere provides system parameters for controlling default platform behavior globally.

System parameters are organized into the following categories:

- *Platform Policy*: Provides platform-level parameters for timeout, management node, deletion, cleanup, concurrency, reconnection, and progress bar policies.
- *Operations Management*: Provides parameters for monitoring, alerts, tasks, and events.
- *Host and VM*: Provides parameters for hosts and VMs.
- *Image Storage*: Provides parameters for images and image storage.
- *Data Storage*: Provides parameters for disks and data storage.
- *Network Resources*: Provides parameters for distributed switches and security groups.

13.6.1 Modify System Parameters

Procedure

1. Navigate to **Menu** > **System Management** > **System Parameters**.
2. Modify the system parameters.

13.6.2 Restore Default System Parameters

Procedure

1. Navigate to **Menu** > **System Management** > **System Parameters**.

2. Click **Reset to Default Settings**.

> **Note:**
> Restoring default settings resets all system parameters, security settings, and advanced HA policy settings to their initial default values. Proceed with caution.

13.6.3 Parameter Entries

13.6.3.1 Platform Policy

Category	Item Name	Item Description
Timeout Policy	Image Addition Timeout	Specify the maximum timeout limit for adding an image. If the image addition process exceeds this time limit, the operation will fail. Default: 3 days. Unit: second, minute, hour, and day.
	Timeout for Creating Disk Image from Disk 1	Specify the maximum timeout limit for creating a disk image from disk 1. If the creation exceeds this time limit , the operation will fail. Default: 3 days. Unit: second, minute, hour, and day.
	VM Creation Timeout	Specify the maximum timeout limit for creating a virtual machine. If the creation exceeds this time limit, the operation will fail. Default: 12 hours. Unit: second, minute , hour, and day.
	VNC Token Timeout	Specify the duration for which a VNC token remains valid . When you launch a VM console, the system generates a token required for the console connection. If you do not reconnect to the VM within the specified timeout period, the token becomes invalid and cannot be used for further console connections. To regain access, simply re -launch the VM console, which will generate a new token . Default: 1,800. Unit: second.
	Platform Service SSH Connection Timeout	Default: 30. Unit: seconds. Specifies the connection timeout for executing SSH commands on platform service VMs.
	WebShell Session Timeout	Specify the duration for which a WebShell session remains valid. Default: 1,800. Unit: second.
	Timeout for Creating Disk Image from Data Disk	Specify the maximum timeout limit for creating a disk image from a data disk. If the creation exceeds the time limit, the operation will fail. Default: 3 days. Unit: second, minute, hour, and day.

Category	Item Name	Item Description
	Timeout for Creating Disk from Disk Image	Specify the maximum timeout limit for creating a disk from a disk image. If the creation exceeds this time limit , the operation will fail. Default: 3 days. Unit: second, minute, hour, and day.
	Max Failed Attempts for Compute Node Inspection	Specify the maximum number of failed attempts allowed when the management node inspects compute nodes. Default: 3.
Management Node Policy	Max User Session Connections	Specify the maximum number of sessions that a single user can start with the management node. If the number of sessions for a user exceeds this limit, new sessions will be denied. Default: 500.
	Image Storage Ping Interval	Specify the time interval at which the management node checks the connectivity of the image storage. If the check (ping) is successful, it indicates that the image storage is in a connected state. Default: 60. Unit: second.
	Orphan GC Scan Interval	Specify the time interval at which the management node scans for orphaned GC (Garbage Collection) jobs. If orphaned GC jobs are detected, the management node takes over the jobs. Default: 60. Unit: second.
	Auto Reconnect Host on Ping Failure	Specify whether to automatically reconnect hosts to the management node, if the management node fails to ping hosts. Default: enabled.
	Host Ping Interval	Specify the time interval at which the management node checks the connectivity of the host. If the check (ping) is successful, it indicates that the host is in a connected state. Default: 60. Unit: second.
	Heartbeat Inspection Interval	Specify the time interval at which the management node writes its heartbeat pulses to the database. Default: 5. Unit: second.
	Data Storage Ping Interval	Specify the time interval at which the management node checks the connectivity of the data storage. Default: 60. Unit: second.
	Check Time Zone in REST API Calls	Specify whether to enable the time zone check when making REST API calls. Deafult: disabled. If enabled, the system will recognize and validate the time zone during API requests. If disabled, the system will only compare local times without considering time zones.

Category	Item Name	Item Description
Deletion Policy	Image Deletion Policy	Specify the deletion policy for images. Default: Delayed Deletion. Options include Delayed Deletion, Immediate Deletion, and Never Delete. • Immediate Deletion: Directly delete the image from the database and file system. • Delayed Deletion: Mark the image as Deleted. The image will be completely deleted from the database and file system after the retention period expires or when you manually force delete the image. The default retention period is 86,400 seconds. • Never Delete: Remove the image from the database but retain its files indefinitely.
	VM Deletion Policy	Specify whether to completely delete the virtual machine or retain it in the recycle bin when you delete a virtual machine. Default: Delayed Deletion. Options include Immediate Deletion, Delayed Deletion, and Never Delete. **1.** Immediate Deletion: • The virtual machine is directly expunged as soon as you delete it. An expunged VM cannot be recovered. Proceed with caution. • Any virtual machines currently in the Recycle Bin will also be expunged as soon as you apply this policy. Proceed with caution. **2.** Delayed Deletion: • The virtual machine is moved to the Recycle Bin after deletion. You can specify a retention period during which the virtual machine is retained in Recycle Bin. After the retention period expires, the virtual machine is automatically expunged and cannot be recovered. • You can recover a deleted virtual machine from the Recycle Bin before it is automatically expunged. **3.** Never Delete: • The virtual machine is moved to the Recycle Bin after deletion and will not be automatically expunged. You can recover a deleted virtual machine from the Recycle Bin.

Category	Item Name	Item Description
		• Virtual machines in the Recycle Bin can be manually expunged. Once expunged, they cannot be recovered. Proceed with caution.
	Disk Deletion Policy	Specify the deletion policy for disks. Default: Delayed Deletion. Options include Immediate Deletion, Delayed Deletion, and Never Delete. When selecting Delayed Deletion, the default retention period is 86,400 seconds. Unit: second, minute, hour, and day.
	Bare Metal Instance Deletion Policy	Specify the deletion policy for bare metal instances . Default: Never Delete. Options include Immediate Deletion and Never Delete. • Immediate Deletion: The bare metal instance is directly expunged as soon as you delete it. • Never Delete: ◦ The bare metal instance is moved to the Recycle Bin after deletion and will not be automatically expunged. You can recover a deleted bare metal instance from the Recycle Bin. ◦ Bare metal instances in the Recycle Bin can be manually expunged. Once expunged, they cannot be recovered. Proceed with caution.
	Sensitive Operation Confirmation Validation	Specify whether to require a secondary confirmation for sensitive operations on critical resources. For example, If you delete a virtual machine, the operation requires an input of specific characters to continue. Default: enabled. If disabled, no additional confirmation is required.
Cleanup Policy	Recycle Bin Expired Resources Cleanup Interval	Specify the time interval for automatic cleanup of expired resources in the recycle bin. Default: 1 hour. Unit: hour. The interval can be set as an integer between 1 and 24. When resources under the delayed deletion policy reach their preset retention time in the recycle bin, the system will check for expired resources at the specified interval and automatically clean them up.
	Standalone Image Storage Data Cleanup Interval	Specify the time interval for automatically cleaning up data in the standalone image storage. Default: 7 days . Unit: second, minute, hour, and day. If disabled, the automatic data cleanup is disabled.

Category	Item Name	Item Description
	Session Cleanup Interval	Specify the time interval for cleanning up sessions that have timed out. Default: 1 hour. Unit: second, minute, hour, and day.
	Data Storage Garbage Collection	Specify whether to enable automatic garbage collection on the data storage after failed attempts to create virtual machines or disks. Default: enabled.
	Data Storage Garbage Collection Interval	Specify the time interval for garbage collection on the data storage. Default: 600. Unit: second.
	Data Storage Garbage Collection Frequency	Specify the number of times the garbage collection will be repeated on the data storage. Default: 50.
	Progress Bar Record Cleanup Interval	Specify the time interval for automatically cleaning up expired progress bar records. Default: 300. Unit: second.
	Failed Disk /Snapshot Cleanup Interval	Specify the time interval for automatically cleaning up failed or invalid disks or snapshots that remain in the SAN storage after deletion attempts. Default: 3,600. Unit : second.
Concurrency Policy	Max Concurrent VM Creations	Specify the maximum number of virtual machines that can be created concurrently on a single host. Default: 1.
	Max Concurrent VM Migrations	Specify the maximum number of virtual machines that can be migrated concurrently when a host is in maintenance mode. Default: 2.
	Image Storage Ping Max Concurrency	Specify the maximum number of image storage checks that can be performed concurrently by the management node. Default: 50.
	Host Allocation Concurrency	Specify whether the host allocation process supports concurrent operations. Default: disabled. If enabled, you can specify the maximum number of host allocators that can be concurrently running. Default: 10. Valid values: 1 to 255. Setting a high concurrent allocation may lead to increased host load.
	Host Concurrenc y Level	Specify the the maximum number of concurrent commands that can be executed on the host. Default: 10

Category	Item Name	Item Description
		. Setting a high concurrency level may lead to increased host load.
	Host Reconnecti on Concurrency	Specify the maximum number of concurrent host reconnection processes that can be initiated by the management node. Default: 100.
	Host Ping Max Concurrency	Specify the maximum number of host connnectivity checks that can be performed concurrently by the management node. Default: 100.
	Concurrent Snapshot Creations on a Single Host	Specify the maximum number of concurrent snapshot creation processes that can be executed on a single host . Default: 10. Setting a high concurrency may lead to increased host load.
	Concurrent Image Distribution to Local Storage	Specify the maximum number of images that can be concurrently distributed to local storage after adding an image. Default: 2.
	Data Storage Ping Max Concurrency	Specify the maximum number of data storage checks that can be performed concurrently by the management node. Default: 50.
Reconnection Policy	Auto Reconnect Distributed Storage Mon Node	Specify whether to automatically reconnect the monitoring nodes of a distributed storage to the management node upon check failures. Default: enabled.
	Distributed Storage Mon Node Reconnecti on Interval	Specify the time interval before the management node of a distributed storage system attempts to reconnect after a failed check. Default: 30. Unit: second.
	Host Max Reconnection Failures	Specify the maximum number of consecutive failed reconnection attempts allowed by the management node before it stops trying to reconnect to a host. Default: 0.
	Auto Reconnect Host	Specify whether to automatically reconnect all hosts while the management node starts its services. Default: enabled.
Progress Bar	Enable LongJob Task Progress Tracking	Enabled by default. When enabled, the platform records the progress of LongJob tasks for display in the task list and operation history. When disabled, the platform no longer records progress data for new LongJob tasks. This setting does not affect LongJob task execution or control the display of other progress bars.

Category	Item Name	Item Description
	Progress Bar Data Retention Period	Specify the maximum duration for which progress bar data is retained in the database. Default: 86,400 seconds . Unit: second, minute, hour, and day.

13.6.3.2 Operations Management

Category	Item Name	Item Description
Monitoring and Alarm	Management Node Log Retention Size	Specify whether to set the maximum size of locally retained management node logs. Logs exceeding this size will be permanently deleted. Default: disabled. Unit: GB. If disabled, there will be no limit on the log retention size.
	Management Node Log Retention Period	Specify whether to set the maximum duration for which management node logs are retained locally. Logs older than this duration will be permanently deleted. Default: disabled. Unit: day. If disabled, there will be no limit on the log retention period.
	Monitoring Data Retention Period	Specify the maximum duration for locally retained monitoring data. Default: 6. Unit: month. Valid values: 1 to 12.
	Monitoring Data Retention Size	Specify the maximum size of locally retained monitoring data. Default: 64GB. Unit: MB, GB, and TB. You can enter an integer that is greater than 1 and is the power of 2.
	Operation Task Retention Period	Specify the maximum duration for retained operation tasks. Operation tasks older than this duration will be permanently deleted. Default: 90. Unit: day.
	Management Node Directory Monitored by System Alarm	Specify the management node directory monitored by the system alarm.
	Auto-Scheduling Task Retention Period	Specify the maximum duration for retained scheduling tasks. Scheduling tasks older than this duration will be permanently deleted. default: 90. Unit: day.
	Event Retention Period	Specify whether to set the maximum duration for retained events. Default: unlimited. If enabled, you can set a retention period. The default retention period is 180 days

Category	Item Name	Item Description
		. Events older than this duration will be permanently deleted.
	Trigger Recovery Check Interval	Specify the time interval at which the system periodically checks whether a trigger has recovered. Default: 10. Unit : second.
	WAL Refresh Interval	Specify the time interval of refreshing Write Ahead Logs (WAL). Default: 2 hours. Unit: minute, hour, and day. If the monitoring data size is too large, it is recommended to set this interval to 30 minutes.
	Alarm Suppressio n Interval	Specify the time interval during which the alarm will not be triggered again after the initial alert. Default: 1,800. Unit: second.
	Alarm Message Retention Period in Cache	Specify the maximum duration for which the total count of alarm messages triggered by resource and event alarms is retained in the cache. Default: 10. Unit: second . Within the duration, user queries for the total count of alarm messages will retrieve the cached value directly. If caching is not required, you can set the value to 0.
	Alarm Inspection Interval	Specify the time interval at which an alarm inspects metric items. Default: 10. Unit: second.
	Alarm Inspection Concurrency	Specify the concurrency that an alarm inspects metric items. Default: 5.
	Alarm Message Threshold for Caching	Specify the threshold for the total count of alarm messages triggered by resource and event alarms that can be written to the cache. Default: 150. If the total count of alarm messages equals or exceeds this threshold, the alarm messages are allowed to be cached . If the count is below the threshold, user queries for the total count of alarm messages will retrieve the real-time value directly.

13.6.3.3 Host and VM

Category	Item Name	Item Description
VM	Overcommit VM Memory	Specify whether to enable memory overcommitment for virtual machines. Default: enabled. If enabled, when assigning a virtual machine to a host, any host where the

Category	Item Name	Item Description
		requested memory quota exceeds the available memory will be filtered out.
	Max Disks per VM	Specify the maximum number of data disks that can be attached to a virtual machine. Default: 24.
	Max CD/DVD Drives per VM	Specify the maximum number of CD/DVD drives that a virtual machine can create. Default: 3. Valid values: 1, 2, and 3. Note that this setting does not take effect on virtual machines that already have more virtual CD/DVD drives than the specified limit.
	AgentPing Interval	Specify the time interval for checking the status of the agent program on the console. Default: 60. Unit: second.
	Agent Idle Timeout	Specify the maximum idle time for the console agent before it times out. Default: 60. Unit: second.
	Inject DNS into VM	Specify whether to inject DNS server addresses into virtual machines. Default: disabled. If enabled, virtual machines will be able to resolve each other's hostnames.
	Max VMs in UI Batch Creation	Specify the maximum number of virtual machines that can be created in a single batch operation via the UI. Default: 100. Valid values: 1 to 10,000.
	User Data Service	Specify whether to enable the User Data service. Default : enabled. If disabled, virtual machines without Userdata settings will not be able to upload internal monitoring data through their agents.
	Display VNC Banner	Specify whether to display a banner notification for VMTools installation at the top of the VNC window. Default: enabled. If disabled, the banner notification will not be shown.
	VM Metadata	Default: Disabled. When enabled, the system creates an independent metadata file for each virtual machine. This feature is currently used only for VM registration. If VM registration is not required, it is recommended to keep this feature disabled.
Host	Host Allocation Strategy	1. Specify the host allocation strategy to be used when creating new virtual machines. Default: Host with min. running VMs.

Category	Item Name	Item Description
		• Host with min. running VMs: The host with the minimum number of running virtual machines will be chosen to create virtual machines. • Host with min. CPU utilization: The host with the minimum CPU utilization will be chosen to create virtual machines. • Host with min. memory utilization: The host with the minimum memory utilization will be chosen to create virtual machines. • Host with max. running VMs: The host with the maximum number of running virtual machines will be chosen to create virtual machines. To use this option, you need to set the maximum number of virtual machines that can run on a host . Then, the system chooses the host that meets the requirements to create virtual machines. If no host is available, you would fail to create a virtual machine. • Host where the VM located last time: When you restart a stopped virtual machine, the system chooses the host where the VM was running last time. If you start a new virtual machine for the first time, the system chooses a host randomly. • Random allocation: The system randomly chooses a host to create virtual machines. **2.** If you choose the "Host with min. CPU utilization" or " Host with min. memory utilization" strategy, you can set the interval for collecting the CPU utilization and memory utilization in the system parameters.
	VM Migration Failure Policy in Maintenance Mode	Specify the policy for handling virtual machine migration failures when host enters maintenance mode. Default: Force Stop VM. • Non-Force Stop VM: The host cannot enter maintenance mode if a virtual machine fails to be migrated while the host is entering maintenance mode . • Force Stop VM: The virtual machine will be forcefully shut down if migration fails while the host is entering maintenance mode.

Category	Item Name	Item Description
	Ignore Errors in Maintenance Mode	Specify whether to allow hosts to ignore errors and report requests as successful while the host is in maintenance mode. Default: disabled.
	Host Upgrade Parallelism	Specify the maximum number of hosts within a cluster that can undergo system upgrades simultaneously. Default: 2.
	Restart Hung KVM Agent	Specify whether to restart the KVM agent when it is detected hung. Default: enabled.
	Sync VM Status on Pinging Host	Specify whether to synchronize the virtual machine status while pinging the host. Default: enabled.
	Host CPU Utilization Collection Interval	Specify the time interval at which the CPU utilization of hosts is collected. The system uses this data to select the host with the lowest CPU utilization for creating new virtual machines. Default: 600. Unit: second.
	Host Memory Utilization Collection Interval	Specify the time interval at which the memory utilization of hosts is collected. The system uses this data to select the host with the lowest memory utilization for creating new virtual machines. Default: 600. Unit: second.
	Host SSH Reconnection Timeout	Specify the timeout period for SSH reconnection attempts to a host. If the SSH connection is not established within this period, the reconnection attempt is considered failed. Default: 300. Unit: second.
	Quick Host Connection Status Detection	Specify whether to enable or disable the quick detection of host connection status. Default: disabled. If enabled, the system will reduce the time interval between connection status checks.

13.6.3.4 Image Storage

Category	Item Name	Item Description
Image	Distribute Image	Specify whether to distribute newly added images to all hosts within the cluster to which the local storage is attached. Default: enabled.
	Auto Delete Temporary Image	Specify whether to automatically delete temporary images generated while cloning virtual machines. Default: enabled.

Category	Item Name	Item Description
Image Storage	Auto Reconnect Image Storage Mon Node	Specify whether to automatically reconnect the monitoring nodes of a distributed image storage to the management node upon check failures. Default: enabled.
	Image Storage Mon Node Reconnection Interval	Default: 30. Unit: Second. The interval that the management node reconnects the monitoring node of the distributed image storage if the monitoring node fails to be detected.

13.6.3.5 Data Storage

Category	Item Name	Item Description
Disk	Local Storage Disk Migration Timeout	Specify the maximum timeout limit for migrating disks from local storage. If the migration process exceeds this time limit, the operation will be considered failed. Default : 1 day. Unit: second, minute, hour, and day.
	Physical I/O Block Size	Specify the size of the physical I/O block that a disk reports to the Hypervisor. Default: 0. Unit: byte. Valid values: 0, 512, and 4096. **Note:** • The value 0 indicates that the size of physical I/O block that a disk takes up is not limited. • This parameter only takes effect on nSDS distributed storage.
	Disk Capacity Refresh Interval	Specify the time interval at which the available space on disks is updated. Default: 3,600. Unit: second. The minimum interval cannot be set lower than 600 seconds.
	Disk Capacity Refresh Scope	Specify the scope of disks that will automatically refresh their capacity information. After setting, the correspond ing disks will automatically refresh their capacity based on the configured refresh interval. Default: All Disks. • All Disks: All disks will be automatically refreshed based on the interval. • Disks with Alarms: Only those disks with alarms will be automatically refreshed based on the interval. • No Refresh: No disks will be automatically refreshed.

Category	Item Name	Item Description
Data Storage	Data Storage Utilization Threshold	Specify a threshold for data storage utilization to prevent overuse of storage space. This is particularly important when storage overcommitment is enabled, as excessive allocation can lead to storage overflow and cause virtual machine storage failures. Default: 0.9. **Note:** When the data storage usage reaches the set threshold, the system will prohibit adding new disks. Existing disks are not affected and work as expected.
	NFS Data Storage GC Cleanup Time Interval	Defaulting to 3600 seconds, used for setting the time interval of NFS data storage GC cleanup and disk snapshot.
	Data Storage Reserved Capacity	Specify a reserved space for data storage. Default: 1GB. Unit: KB, MB, GB, and TB. • For local storage and distributed storage, note that the system will set this reserved capacity for each storage pool (i.e., the actual effective reserved capacity = the specified reserved capacity value × number of storage pools). Please set a reasonable value to avoid exceeding the total physical data storage capacity. • For other storage, the reserved capacity is set to the specified value. **Note:** The reserved capacity of data storage does not affect the normal read/write operations of existing disks.
	Distributed Storage GC Interval	Specify the time interval at which the distributed storage system performs garbage collection to clean up unused or orphaned data on disks. Default: 3,600. Unit: second.
	nSDS Distributed Storage Image Cache Cleanup Interval	The interval of cleaning image caches from nSDS distributed storage. This setting is applied after images in distributed storage are deleted and image caches are no longer used by virtual machines. Default: 43,200. Unit: second.

Category	Item Name	Item Description
	Distributed Storage Pool Deletion Policy	Specify whether to delete all storage pools when deleting distributed storage. Default: disabled. If enabled, this may result in data loss. Proceed with caution.
	Non-Distributed Storage Image Cache Cleanup Interval	The interval of cleaning image cache from non-distributed storage. This setting is applied after images in data storage (non-distributed storage) are deleted and image caches are no longer used by any virtual machines. Default: 86,400. Unit: second.
	Disable queue_if_no_path for Multipath Device	Specify whether to disable the queue_if_no_path feature for multipath devices. Default: enabled. The system will return I/O failures if all paths to the multipath device used by SAN storage are lost. If disabled, it may impact the reliability of heartbeat detection and failover mechanisms in the platform. Proceed with caution.
	Initial Size for Thin-Provisioned Disk	Specify the initial size of a thin-provisioned disk. Default: 5GB. Unit: MB, GB, and TB. The initial size cannot be lower than 1 GB.
	Min Available Space for Thin-Provisioned Disk	Specify the minimum available space threshold for a thin-provisioned disk. When the difference between the actual size of a thin-provisioned disk and the used storage space falls below this threshold, the disk will automatically expand based on a predefined size. Default: 5GB. Unit: MB, GB, and TB.
	Auto Expansion Size for Thin-Provisioned Disk	Specify the size by which a thin-provisioned disk will automatically expand. Default: 5GB. Default: MB, GB, and TB.
	Max Usage Ratio for Thin-Provisioned Disk	Specify the maximum usage ratio for a thin-provisioned disk. When the actual usage ratio of a thin-provisioned disk exceeds the maximum usage ratio, the disk will automatically expand based on a predefined size. Default: 85. Unit: %.
	Multipath Access to Data Storage	Specify the control policy of multipath access to data storage. • Enable: Enable the multipath service of the platform. This service is managed by the management node, and thus you cannot customize the multipath service. • Disable: Disable the multipath service of the platform. If disabled, only single path access is supported.

Category	Item Name	Item Description
		• Custom: Enable the multipath service of the platform. In addition, you can customize the multipath service. If you customize the service, the customization prevails over the platform configuration managed by the management node. **Note:** If you modify this setting, you need to reconnect a host to make the modification take effect on the host.

13.6.3.6 Network Resources

Category	Item Name	Item Description
Distributed Switch	Distributed Switch NoVLAN Network MTU	Specify the maximum packet size that can be transmitted over a NoVLAN network. Default: 1,500. Unit: byte.
	Distributed Switch VLAN Network MTU	Specify the maximum packet size that can be transmitted over a VLAN network. Default: 1,500. Unit: byte.
Security Group	Default Egress Policy of Security Group	Specify the default egress policy of security groups. Default: accept. Valid values: accept and deny.
	Default Ingress Policy of Security Group	Specify the default ingress policy of security groups. Default: drop. Valid values: accept, drop, and deny.

Glossary

14.1.1.1 Glossary

Data Center
A data center is the largest resource namespace within a virtualization platform, including resources such as clusters, hosts, data storage, distributed switches, and distributed port groups.

Cluster
A logical collection of a group of hosts (compute nodes).

Host
A host is an x86 or ARM physical server running a KVM virtualization hypervisor, providing resources such as computing, networking, and storage to virtual machines.

Virtual Machine Group
A logical grouping of virtual machines based on business needs.

Virtual Machine
A virtual machine is a virtualized host running on a physical host, capable of running an operation system and applications just like a physical host.

Disk
A disk provides storage space for a virtual machine. Disks are categorized into system disks and data disks.

System Disk
A system disk provides support for the system operations of a virtual machine.

Data Disk
A data disk provides extended storage space for a virtual machine.

Image
An image is a template file used by virtual machines or disks. Images are categorized into system images and disk images.

Image Storage
An image storage is a virtualized resource that provides storage space for image template files used by virtual machines or disks. An image storage can be categorized into standalone image storage and distributed image storage.

Standalone Image Storage
A standalone image storage stores image files through image slices and support incremental storage.

Distributed Image Storage
A distributed image storage stores image files through distributed block storage.

Data Storage
A data storage is a virtualized resource that provides storage space for virtual machines and their application data. A data storage can be categorized into local storage and network shared storage .

Local Storage
A local storage is storage resource constructed using the physical storage space of one or more hosts.

Netowrk Shared Storage
A storage system used for remote storage of virtual machines and their application data, accessible concurrently by hosts over a network.

Distributed Switch
A virtual switching device that provides unified virtual network management and monitoring for virtual machines within a cluster.

Distributed Port Group
A logical grouping of ports on a distributed switch, used for port configuration.

vNUMA Configuration

vNUMA uses CPU pinning to passthrough the topology of associated host physical NUMA (pNUMA) nodes to a virtual machine, generating a topology of virtual NUMA (vNUMA) nodes for the virtual machine. This topology enables a vCPU on a vNUMA node to primarily access the local memory and thus improves VM performance.

NUMA (Non-Uniform Memory Access)

Non-uniform memory access (NUMA) is a computer memory design where the memory access time depends on the memory location relative to the CPU. Under NUMA, a processor can access its own local memory faster than non-local memory and thus improves VM performance.

pNUMA Node (physical NUMA Node)

A pNUMA node (physical NUMA node) is a host NUMA node predefined based on the host NUMA architecture. It is used to manage the CPUs and memory of the host.

pNUMA Topology (physical NUMA Topology)

A pNUMA topology (physical NUMA topology) is the topology of the host NUMA nodes predefined by the CPU vendor based on the host NUMA architecture.

vNUMA Node (virtual NUMA Node)

A vNUMA node (virtual NUMA node) is generated by passing-through associated pNUMA nodes via CPU pinning. It is used to manage the CPUs and memory of a virtual machine.

vNUMA Topology (virtual NUMA Topology)

A vNUMA topology (virtual NUMA topology) is the topology of VM NUMA nodes generated by passing-through associated pNUMA nodes via CPU pinning.

Local Memory

Local memory is the memory that a CPU (pCPU or vCPU) accesses through the Uncore iMC (Integrated Memory Controller) of the same NUMA (pNUMA or vNUMA) node. Compared with accessing non-local memory, accessing local memory has lower latencies.

CPU Pinning

CPU pinning assigns the virtual CPUs (vCPUs) of a virtual machine to specific physical CPUs (pCPUs) of the host, which improves VM performance.

EmulatorPin Configuration

EmulatorPin assigns all other threads than virtual CPU (vCPU) threads and IO threads of a virtual machine to physical CPUs (pCPUs) of the host so that these threads run on assigned pCPUs.

Snapshot

A snapshot is a point-in-time capture of data status in a disk.

VM Scheduling Policy

A VM scheduling policy is a resource orchestration policy based on which virtual machines are assigned to hosts to achieve the high performance and high availability of businesses.

Management Node

A physical host where the system is installed, providing UI management and virtualization platform deployment capabilities.

Compute Node

Also known as a host, it is a physical server that provides computing, networking, and storage resources for virtual machines.

iSCSI Storage

iSCSI storage is an SAN storage that uses the iSCSI protocol for data transmission. You can add an iSCSI SAN block as a SAN storage or passthrough the block to a virtual machine.

FC Storage

FC storage is an SAN storage that uses the FC technology for data transmission. You can add an FC SAN block as a SAN storage or passthrough the block to a virtual machine.

NVMe Storage

A type of storage implemented via the NVMe-oF (NVMe over fabrics) protocol. You can add a block device configured from an NVMe storage as SAN storage.

Management Network

A management network is used to manage physical resources in the platform. For example, you can create a management network to manage access to hosts, data storage, and image storage.

Security Group

A security group provides security control services for VM NICs. It filters the ingress or egress TCP, UDP, and ICMP packets of VM NICs based on the specified security rules.

MN Monitoring

Management Node (MN) monitoring allows you to view the health status of each management node when you use multiple management nodes to achieve high availability.

Alarm

An alarm is used to monitor the status of time-series data and events and respond to the status change. Alarms can be categorized into resource alarm and event alarm.

Message Template

A message template specifies the text template of a resource alarm message or event alarm message sent to an SNS system.

Endpoint

An endpoint is a method that users obtain subscribed messages.

Alarm Message

An alarm message is a message sent the time when an alarm is triggered.

Operation Task

An operation task is a chronological record of operations on the specified objects and their operation results.

HA Task

HA task logs generated when the platform executes high availability procedures in accordance with the enabled HA policy.

Scheduling Task

Scheduling task logs generated when the platform executes dynamic resource scheduling operations after the cluster DRS is enabled.

Event

Event monitors and records all activities on the platform. You can use this feature to implement operation tracking, security analysis, troubleshooting, and automatic O&M.

Backup Plan

You can create a backup plan to back up local virtual machines or platform databases to a specified backup storage on a regular basis.

Local Backup Data

Local backup data of virtual machines or platform databases is stored in the local backup storage.

Local Backup Storage

A local backup storage is located at the local data center and is used to store local backup data.

Remote Backup Storage

A remote backup storage is located at the a remote data center and is used to store remote backup data.

Tag

A tag is used to mark resources. You can use a tag to search for and aggregate resources.

Custom Attribute

Custom attributes are user-defined key-value pairs in the platform that extend the resource tagging and management capabilities.

Single Sign-On

The Single Sign-On service supports seamless access to SSO systems. Through the service, related users can directly log in to the platform and manage resources.

Console Proxy

Console proxy allows you to log in to a virtual machine by using the IP address of a proxy.

AccessKey Management

An AccessKey pair is a security credential that one party authorizes another party to call API operations and access its resources in the platform. AccessKey pairs shall be kept confidential.

IP Blocklist/Allowlist

An IP blocklist or allowlist identifies and filters IP addresses that access the platform. You can create an IP allowlist or blocklist to improve access control of the platform.

User Management

A user can be created by the admin or synchronized from an authentication system. A user is managed by the admin. Resources created by a user are managed by the user.

Theme and Appearance

You can customize the theme and appearance of the platform.

Email Server

If you select Email as the endpoint of an alarm, you need to set an email server. Then alarm messages are sent to the email server.

Log Server

A log server receives management node logs or platform operation logs sent from the virtualiza tion platform for centralized retention, search, and analysis.

System Parameters

System parameters allow you to configure settings that take effect on the whole platform.

HA Policy

HA Policy is a mechanism that ensures sustained and stable running of the business if virtual machine are unexpectedly stopped or are errored because of errors occurring to compute, network, or storage resources associated with the virtual machines. By enabling this feature, you can customize VM HA policies to ensure your business continuity and stability.

Bare Metal Template

With bare metal templates, preconfigured files can be quickly generated to achieve unattended bulk OS installation for bare metal instances.

Bare Metal Cluster

A bare metal cluster consists of bare metal chassis.

Deployment Server

An independent server used for providing PXE services and console proxies for bare metal chassis.

Bare Metal Chassis

Bare metal chassis is used to create bare metal instances and can be universally identified by the BMC interface and IPMI configuration.

Bare Metal Instance

A virtualized instance of bare metal chassis.