



NexaSphere User Guide

Version: V2.0.2

Issue: V2.0.2

Contents

1 Introduction.....	1
2 About NexaSphere.....	2
3 NexaSphere UI.....	3
4 Product Features.....	5
5 Log in to NexaSphere.....	30
6 License Service.....	31
6.1 Overview.....	31
6.2 Use License.....	32
6.2.1 USB Key Licensing.....	32
6.2.2 Request Key Licensing.....	33
6.2.3 License Status Descriptions.....	34
7 Initialize NexaSphere.....	37
7.1 Initialize Manually.....	37
7.2 Restore from Backup Data.....	40
8 Platform Resources Definition.....	42
8.1 Resource Definition.....	42
8.1.1 Compute Resource Data Center & Cluster & Host.....	42
8.1.2 Compute Resource Image & Image Storage.....	44
8.1.3 Compute Resource Virtual Machine & VM Group.....	44
8.1.4 Storage Resource Data Storage.....	46
8.1.5 Network Resource Distributed Switch & Distributed Port Group.....	47
8.2 Resource Relationship.....	47
8.2.1 Platform Resource Relationship.....	48
8.2.2 Cluster Data Storage Image Storage Relationship.....	50
9 Compute Management.....	52
9.1 Data Center.....	52
9.1.1 Data Center Basic Operations.....	52
9.2 Cluster.....	53
9.2.1 Cluster Basic Operations.....	53
9.2.2 Cluster DRS.....	56
9.3 Host.....	58
9.3.1 Host Basic Operations.....	58
9.3.2 Host Hardware Device.....	61
9.3.2.1 Host NUMA Topology.....	61
9.3.2.2 LUN.....	63
9.3.2.3 Host NIC.....	63
9.3.2.4 GPU Device.....	67
9.3.2.5 USB Device.....	68
9.3.2.6 PCIe Device.....	69

9.4 Image Storage.....	70
9.4.1 Image Storage Basic Operations.....	70
9.4.1.1 Add a Standalone Image Storage.....	70
9.4.1.2 Add a Distributed Image Storage.....	72
9.4.1.3 Manage an Image Storage.....	73
9.4.2 Image Basic Operations.....	73
9.4.2.1 Add an Image.....	73
9.4.2.2 Synchronize Images.....	74
9.4.2.3 Migrate an Image.....	75
9.4.2.4 Delete an Image.....	75
9.4.3 Image Cross-Platform Usage.....	76
9.4.3.1 Export Image.....	76
9.4.3.2 Export VM as OVA Template.....	76
9.5 Virtual Machine.....	77
9.5.1 Virtual Machine Basic Operations.....	77
9.5.1.1 Create a New Virtual Machine.....	77
9.5.1.2 Import a Virtual Machine.....	85
9.5.1.3 Create a Virtual Machine from a Template.....	87
9.5.1.4 Register a Virtual Machine.....	89
9.5.1.5 Clone a Virtual Machine.....	90
9.5.1.6 Access a Virtual Machine.....	92
9.5.1.6.1 (Optional) Manage VM Access and Boot Options.....	92
9.5.1.6.2 Access a VM by using Console.....	95
9.5.1.6.3 Access a VM by using SSH.....	96
9.5.1.7 Modify a Virtual Machine.....	96
9.5.1.8 Delete a Virtual Machine.....	98
9.5.2 VM Group Management.....	99
9.5.3 Virtual Machine VMTools.....	100
9.5.3.1 Install VMTools on Linux VMs.....	103
9.5.3.2 Install VMTools on Windows VMs.....	104
9.5.3.3 VMTools Compatibility with OS.....	104
9.5.4 VM Failure Management.....	108
9.5.5 VM Time Synchronization.....	108
9.5.6 VM User Data.....	110
9.5.6.1 Import User Data to Linux VMs.....	110
9.5.6.2 Import User Data to Windows VMs.....	111
9.5.7 VM Template.....	112
9.5.7.1 Clone a Virtual Machine to a Template.....	112
9.5.7.2 Convert a Virtual Machine to a Template.....	113
9.5.7.3 Convert a Template to a Virtual Machine.....	114
9.5.7.4 Create a Virtual Machine from a Template.....	114
9.5.7.5 Manage Virtual Machine Templates.....	117
9.5.8 VM Specification.....	118
9.5.8.1 Create a Windows VM Specification.....	118

9.5.8.2 Create a Linux VM Specification.....	120
9.5.8.3 Manage VM Specifications.....	121
9.5.9 VM Migration Management.....	121
9.5.9.1 Change Host.....	121
9.5.9.1.1 Hot Migration.....	121
9.5.9.1.2 Cold Migration.....	123
9.5.9.2 Change Data Storage.....	123
9.5.9.2.1 Hot Migration from SAN Storage to SAN Storage.....	123
9.5.9.2.2 Cold Migration from SAN Storage to SAN Storage.....	124
9.5.9.2.3 Hot Migration Between SAN Storage and nSDS Distributed Storage.....	125
9.5.9.3 Change Host and Data Storage.....	126
9.5.9.3.1 Hot Migration Across Data Storage of the Same Type.....	126
9.5.9.3.2 Hot Migration Across Data Storage of Different Types.....	127
9.5.9.3.3 Hot Migration Across Storage Pools Within the Same nSDS Distributed Storage.....	128
9.5.9.3.4 Cold Migration Across Data Storage of the Same Type.....	129
9.5.9.3.5 Cold Migration Across Storage Pools Within the Same nSDS Distributed Storage.....	130
9.5.10 VM Performance.....	131
9.5.10.1 VM Resource Contention.....	131
9.5.10.2 VM CPU.....	133
9.5.10.3 VM QoS.....	134
9.5.11 VM Snapshot Management.....	136
9.5.11.1 Overview.....	136
9.5.11.2 Snapshot Basic Operations.....	136
9.5.11.2.1 Create a Snapshot.....	136
9.5.11.2.2 Revert to a Snapshot.....	137
9.5.11.2.3 New Virtual Machine from Snapshot.....	138
9.5.11.2.4 View Snapshot.....	139
9.5.11.2.5 Delete a Snapshot.....	140
9.5.11.3 Snapshot Policy Basic Operations.....	140
9.5.11.3.1 New Snapshot Policy.....	140
9.5.11.3.2 Enable/Disable Snapshot Policy.....	141
9.5.11.3.3 Modify Configuration.....	141
9.5.11.3.4 Delete Snapshot Policy.....	142
9.5.11.4 Snapshots Usage Recommendations.....	142
9.5.12 VM Scheduling Policy.....	143
9.5.12.1 Overview.....	143
9.5.12.2 Create Mutually Exclusive/Affinitive VM Scheduling Policies.....	151
9.5.12.3 Create VM Mutually Exclusive/Affinitive Host Scheduling Policies.....	152
9.5.12.4 Manage VM Scheduling Policies and Related Resources.....	154
9.5.13 VM HA.....	156
9.5.13.1 Overview.....	156
9.5.13.2 HA Policy Basic Operations.....	159

9.5.13.3 Implement HA Policy in Business Practices.....	163
9.5.14 VM Security.....	164
9.5.14.1 Configure and Manage a Standard Key Provider.....	164
9.5.14.1.1 Add a Standard Key Provider.....	164
9.5.14.1.2 Establish Trust.....	165
9.5.14.1.2.1 Make KMS Trust NexaSphere.....	165
9.5.14.1.2.2 Make NexaSphere Trust the KMS.....	166
9.5.14.1.3 Delete a Standard Key Provider.....	167
9.5.14.1.4 Use the Thales Key Provider.....	167
9.5.14.1.5 Use the HashiCorp Key Provider.....	178
9.5.14.2 Configure and Manage a Native Key Provider.....	182
9.5.14.2.1 Add a Native Key Provider.....	182
9.5.14.2.2 Back Up the Native Key Provider.....	183
9.5.14.2.3 Restore the Native Key Provider.....	184
9.5.14.2.4 Delete the Native Key Provider.....	184
9.5.14.3 Rekey.....	185
9.5.14.4 Secure VMs with vTPM.....	185
10 Storage Management.....	187
10.1 Add a Data Storage.....	187
10.1.1 Add a Local Storage.....	187
10.1.2 Add an NFS Storage.....	188
10.1.3 Add a SAN Storage.....	189
10.1.3.1 Add an iSCSI Storage.....	191
10.1.3.2 Synchronize a FC Storage.....	191
10.1.3.3 Add a NVMe Storage.....	192
10.1.4 Add a nSDS Distributed Storage.....	193
10.1.5 Add a ZHPS Distributed Storage.....	194
10.1.6 Add a ZBS Distributed Storage.....	195
10.2 Modify Data Storage Configuration.....	196
10.2.1 Modify Local Storage Configuration.....	196
10.2.2 Modify NFS Storage Configuration.....	197
10.2.3 Modify SAN Storage Configuration.....	198
10.2.4 Modify Distributed Storage Configuration.....	200
10.3 Data Storage Cleanup & Deletion.....	202
11 Network Management.....	204
11.1 Network Resource.....	204
11.1.1 Distributed Switch.....	204
11.1.1.1 Create a Distributed Switch.....	204
11.1.1.2 Distributed Switch Uplink.....	208
11.1.1.2.1 Modify Uplink Configurations.....	208
11.1.1.2.2 Manage Joined Hosts.....	208
11.1.1.2.3 Configure Uplinks for Hosts.....	209
11.1.1.3 Network Topology.....	210

11.1.1.3.1 View Network Topology.....	210
11.1.1.3.2 Supported Network Topology Operations.....	210
11.1.1.4 Attach/Detach a Distributed Switch to/from a Cluster.....	210
11.1.1.5 Delete a Distributed Switch.....	211
11.1.2 Distributed Port Group.....	212
11.1.2.1 Create a Distributed Port Group.....	215
11.1.2.2 Modify Distributed Port Group Configurations.....	217
11.1.2.3 Delete a Distributed Port Group.....	217
11.1.3 Kernel Adapter.....	218
11.1.3.1 Create a Kernel Adapter.....	218
11.1.3.2 Modify Kernel Adapter Configurations.....	218
11.1.3.3 Delete a Kernel Adapter.....	219
11.2 Network Service.....	219
11.2.1 Security Group.....	219
11.2.1.1 Overview.....	219
11.2.1.2 Create a Security Group and Related Rules.....	221
11.2.1.3 Modify Security Groups and Related Rules.....	224
12 O&M Management.....	226
12.1 Resource Monitoring.....	226
12.1.1 Resource Performance Monitoring.....	226
12.1.1.1 Overview.....	226
12.1.1.2 Capacity Monitoring.....	226
12.1.1.3 View Monitoring Charts.....	230
12.1.1.4 Customize Monitoring Charts.....	230
12.1.2 Dashboard Monitoring.....	234
12.1.3 Dual Management Node Monitoring.....	235
12.1.4 Host Hardware Monitoring.....	236
12.2 Alarm Service.....	236
12.2.1 Overview.....	236
12.2.2 Endpoint.....	237
12.2.2.1 New Endpoint.....	237
12.2.2.1.1 Email.....	237
12.2.2.1.2 DingTalk.....	238
12.2.2.1.3 Lark.....	239
12.2.2.1.4 WeCom.....	240
12.2.2.1.5 SMS.....	241
12.2.2.1.6 HTTP Application.....	242
12.2.2.1.7 Microsoft Teams.....	243
12.2.2.1.8 SNMP Trap Receiver.....	243
12.2.2.2 Manage Endpoint.....	244
12.2.3 Message Template.....	245
12.2.3.1 New Message Template.....	245
12.2.3.2 Manage Message Template.....	246
12.2.4 Alarm.....	247

12.2.4.1 Alarm Rules.....	247
12.2.4.2 New Resource Alarm.....	248
12.2.4.3 New Event Alarm.....	249
12.2.4.4 Manage Alarm.....	250
12.2.5 Alarm Message.....	251
12.2.5.1 View Alarm Messages.....	251
12.2.5.2 Acknowledge Alarm Messages.....	252
12.2.5.3 Set Silence Period for Alarm Messages.....	252
12.2.5.4 Restore Alarms.....	252
12.3 Task.....	253
12.4 Event.....	254
12.5 Log Collection.....	255
12.5.1 Collect Logs.....	255
12.5.2 Manage Collected Logs.....	256
12.6 Tag Management.....	256
12.6.1 Overview.....	256
12.6.2 Create a Tag.....	257
12.6.3 Attach/Detach a Tag.....	258
12.6.4 Delete a Tag.....	258
12.6.5 Search Resources Using Tags.....	258
12.7 Custom Attribute.....	259
12.7.1 Create a Custom Attribute.....	259
12.7.2 Add and Edit Custom Attributes.....	259
12.7.3 Manage a Custom Attribute.....	260
12.7.4 Search Resources Using Custom Attributes.....	260
12.8 Migration Service.....	261
12.8.1 Deploy Migration Service.....	261
12.8.1.1 Step 1: Upload Migration Service Package.....	261
12.8.1.2 Step 2: Deploy Migration Service.....	262
12.8.2 Manage Migration Service.....	263
12.8.2.1 Upgrade the Migration Service.....	263
12.8.2.2 Enable/Disable the Migration Service.....	263
12.8.2.3 Delete the Migration Service.....	264
12.8.3 Manage Migration Resources.....	264
12.8.3.1 Add a Source Platform.....	264
12.8.3.2 Delete a Source Platform.....	265
12.8.3.3 Add a Data Gateway.....	265
12.8.3.4 Delete a Data Gateway.....	266
12.8.4 Manage Migration Tasks.....	267
12.8.4.1 Create a Migration Task.....	267
12.8.4.2 Synchronize a Migration Task.....	269
12.8.4.3 Create a Test Virtual Machine.....	269
12.8.4.4 Perform Cutover.....	270
12.8.4.5 Delete a Migration Task.....	271

13 System Management.....	272
13.1 Identity and Access Management.....	272
13.1.1 Overview.....	272
13.1.2 Preparation.....	272
13.1.3 Single Sign-On.....	272
13.1.3.1 Add OIDC SSO Server.....	273
13.1.3.2 Add AD SSO Server.....	274
13.1.3.3 Add LDAP SSO Server.....	275
13.1.3.4 Manage SSO Server.....	277
13.1.4 Role Management.....	277
13.1.4.1 System Predefined Roles.....	277
13.1.4.2 Create Custom Role.....	278
13.1.4.3 Clone Role.....	278
13.1.4.4 Modify Role Permissions.....	279
13.1.4.5 Delete Role.....	279
13.1.5 User Management.....	279
13.1.5.1 New User.....	280
13.1.5.2 Disable/Enable User.....	281
13.1.5.3 Modify User Configuration.....	281
13.1.5.4 Change User Password.....	281
13.1.5.5 Change a Regular User to an Admin User.....	282
13.1.5.6 Delete a User.....	282
13.1.6 User Group Management.....	282
13.1.6.1 New User Group.....	283
13.1.6.2 Modify User Group Configuration.....	284
13.1.6.3 Delete User Group.....	284
13.2 Security Access Settings.....	284
13.2.1 IP Blocklist and Allowlist Management.....	285
13.2.2 Certificate Management.....	286
13.2.2.1 Import Third-Party Certificate.....	286
13.2.2.2 Import System Self-Signed Certificate.....	287
13.2.2.3 Update Certificate.....	289
13.2.2.4 Switch to HTTP Login.....	289
13.2.3 Security Settings.....	290
13.3 System Settings.....	292
13.3.1 AccessKey Management.....	292
13.3.2 Console Proxy Management.....	293
13.3.3 SNMP Management.....	293
13.3.3.1 Overview.....	293
13.3.3.2 Enable SNMP Management.....	294
13.3.3.3 Modify SNMP Configuration.....	295
13.3.3.4 Download MIB File.....	295
13.3.3.5 Manage SNMP Trap Receiver.....	295
13.3.3.6 Disable SNMP Management.....	296

13.3.4 Time Configuration.....	296
13.3.4.1 Overview.....	296
13.3.4.2 Modify Time Server Configuration.....	298
13.3.4.3 Synchronize Time.....	299
13.3.5 Log Server.....	299
13.3.5.1 Add Log Server.....	299
13.3.6 Email Server.....	300
13.3.6.1 Add Email Server.....	300
13.3.6.2 Manage Email Server.....	300
13.3.7 Theme Appearance.....	300
13.3.7.1 Customize Theme and Appearance.....	300
13.3.7.2 Restore to Default Theme and Appearance.....	301
13.3.8 System Parameters.....	302
14 Backup Management.....	304
14.1 Overview.....	304
14.2 Preparation.....	305
14.3 Add a Backup Storage.....	305
14.3.1 Add a Local Backup Storage.....	305
14.3.1.1 Ruse Image Storage.....	306
14.3.1.2 Ruse Host.....	306
14.3.1.3 Dedicated Backup Storage.....	307
14.3.2 Add Remote Backup Storage.....	309
14.3.2.1 Dedicated Backup Storage.....	309
14.4 Virtual Machine Backup.....	310
14.4.1 Backup Policy.....	310
14.4.2 New Backup Plan.....	312
14.4.3 View Backup Plan and Data.....	314
14.4.3.1 View Backup Plan Executions.....	315
14.4.3.2 View VM Backup Data.....	316
14.4.4 Perform Backups Manually.....	316
14.4.4.1 Immediate Backup.....	316
14.4.4.2 Create On-Demand Backup.....	317
14.5 Platform Database Backup.....	317
14.5.1 Backup Policy.....	318
14.5.2 New Backup Plan.....	319
14.5.3 View Backup Plan and Data.....	320
14.5.3.1 View Backup Plan Executions.....	321
14.5.3.2 View Platform Database Backup Data.....	321
14.6 Data Recovery.....	322
14.6.1 Restore Virtual Machine.....	322
14.6.2 New Virtual Machine from Backup.....	323
14.6.3 Restore Platform Database.....	324
14.7 Backup Data Management.....	325
14.8 Backup Plan Management.....	327

14.9 Backup Storage Management.....	328
15 Bare Metal Management.....	332
15.1 Overview.....	332
15.2 Preparation.....	334
15.3 Quick Start Guide.....	335
15.4 Bare Metal Template.....	336
15.4.1 Template Syntax Rules.....	336
15.4.2 System Preconfigured Templates.....	337
15.4.3 Add a Custom Template.....	338
15.5 Bare Metal Cluster.....	339
15.5.1 New Bare Metal Cluster.....	339
15.5.2 Deployment Server Management.....	340
15.5.3 Delete a Bare Metal Cluster.....	340
15.6 Bare Metal Chassis.....	341
15.6.1 Add a Bare Metal Chassis.....	341
15.6.1.1 Manually Add a Bare Metal Chassis.....	341
15.6.1.2 Import a Bare Metal Chassis from a Template.....	342
15.6.2 Bare Metal Chassis Status Management.....	342
15.6.3 Access a Bare Metal Chassis.....	343
15.6.4 Delete a Bare Metal Chassis.....	344
15.7 Bare Metal Instance.....	344
15.7.1 New Bare Metal Instance.....	344
15.7.2 Bare Metal Instance Status Management.....	345
15.7.3 Delete a Bare Metal Instance.....	346
16 Storage Service.....	347
16.1 Deploy Distributed Storage in 3 Steps.....	347
16.1.1 Step 1: Upload Installation Package.....	347
16.1.2 Step 2: Deploy Management Service.....	348
16.1.3 Step 3: Initialize Distributed Storage.....	348
16.2 Take Over Existing Distributed Storage.....	350
16.2.1 Take Over Distributed Storage.....	350
16.2.2 Cancel Takeover of Distributed Storage.....	350
16.3 Distributed Storage Resource Management.....	351
16.3.1 Storage Pool.....	351
16.3.1.1 Create a General Purpose Pool.....	351
16.3.1.2 Manage a General Purpose Pool.....	355
16.3.2 Storage Node.....	359
16.3.2.1 Add a General Purpose Storage Server.....	359
16.3.2.2 Manage a General Purpose Storage Server.....	365
16.3.3 Data Disk.....	366
16.3.3.1 Create a Data Disk on General Purpose Node.....	366
16.3.3.2 Manage a Data Disk on General Purpose Node.....	367
16.3.4 Physical Hard Disk.....	367

16.3.4.1 Scan Hard Disks on General Purpose Node.....	367
16.3.4.2 Manage Hard Disks on General Purpose Node.....	368
18 Glossary.....	370

1 Introduction

About This Guide

This guide will help you quickly get started with NexaSphere 2.0.2, covering product architecture, features, system configuration, and detailed operational instructions.

Intended Audience


This document is intended for the following readers:

- Technical Support Engineers
- Maintenance Engineers
- Product Consulting Engineers
- Anyone interested in NexaSphere

2 About NexaSphere

NexaSphere is built on the core principles of high performance, high security, high stability, and high self-developed capabilities. It features the cloud platform engine, along with 4S characteristics. By utilizing cutting-edge technologies such as server virtualization, network virtualization, and storage virtualization, it also offers advanced intelligent operation and maintenance capabilities. You can quickly build a virtualized data center with NexaSphere, and combine it with Nexavm Technologies AG's rich product line to create an integrated solution ranging from IaaS to PaaS.


Figure 2-1: Functional Architecture



3 NexaSphere UI

The user interface of NexaSphere is designed to streamline resource management, offering intuitive navigation and intelligent contextual views for efficient resource creation, monitoring, and operations.

Figure 3-1: User Interface



- **Navigation Pane:** Located at the left of the interface, navigation pane provides quick access to platform's core functions, such as dashboard, inventory, system settings. The navigation pane and resource tree form the primary framework for system navigation. You can switch between modules seamlessly through the first and second-level menus in the navigation pane.
- **Resource Tree:** Located under the secondary menu of the inventory, resource tree hierarchically organizes resources, such as virtual machines, data storage, distributed switches, and more. You can expand/collapse nodes for structured resource management and perform right-click actions.
- **Task/Alarms Panel:** Displays platform tasks and triggered alarms at the interface bottom. The task/alarms panel is globally visible, ensuring you can promptly track task execution and changes in platform status. You can expand/collapse the panel.
- **Details Page:** The main workspace of the interface. Resource details page displays comprehensive information about the selected resource, including overview, monitoring, associated resources, and more. Details page may be presented in card, list, or other formats.

- **Action Bar:** Located at the top or side of the details page, action bar provides operation buttons that you can use to perform specific management actions on the current resource.
- **Tabs:** Located near the top of the details page, tabs allow you to switch between different functional views within the same page. Tabs provide access to related resource tabs and, when combined with the action bar, you can achieve a seamless "select-view-operate" workflow.

4 Product Features

NexaSphere provides multiple licenses and the features offered varies with different licenses. This section mainly introduces the feature scopes provided by **Advanced-Paid** license.

Table 4-1: Feature List

Type	Features	Description
Dashboard	Dashboard	Provides multi-dimensional data statistics cards.
		Provides a default homepage customized for each user's perspective.
	Resource Overview	Supports viewing basic information, capacity information, configuration details, hardware details, and related objects in card format.
Root Node	Root Node	Allows you to manage data center resources under the root node.
Data Center	Lifecycle Management	Supports creating data centers on demand and allows setting up independent clusters, storage, and network resources within each data center.
		Allows you to manage the lifecycle of data centers, such as creating and deleting data centers.
	Associated Resources	Allows you to manage the associated clusters , hosts, virtual machines, data storage, image storage, storage target, and network resources in a data center.
	Recycle Bin	Allows you to manage the recycling of virtual machines, images, and disks.
	Export List	Allows you to manage the exported virtual machines and images, including direct download and copying download links.
	Resource Topology	Displays data center resource scales using topology graphics.
Cluster	Lifecycle Management	Allows you to manage the lifecycle of clusters, such as creating, modifying, deleting clusters.
	Monitoring	Allows you to view visualized monitoring data of all hosts and VMs in the cluster, including CPU, memory, disk, and NIC.

Type	Features	Description
	Associated Resources	Allows you to manage the associated hosts, virtual machines, data storage, and network resources in a cluster.
	Dynamic Resource Scheduling	Allows you to monitor the CPU and memory load of host on a cluster basis and dynamically adjust the virtual machine workloads based on scheduling policies.
		Supports manual DRS (provides scheduling suggestions based on which you can schedule resources for load balancing) and auto DRS (schedules resources based on the system scheduling algorithm without arousing your awareness).
	Network Settings	Allows you to specify a VDI network and migration network for a cluster.
	Overcommitment	Allows you to set CPU and memory overcommit ratios for clusters to increase compute resource utilization.
	Host Settings in Cluster	Allows you to set default parameters for hosts on a cluster basis, including Host CPU Model Check, ignore_msrs Option, Host Zero Copy, Huge Pages, and Host Reserved Memory.
Host	VM Settings in Cluster	Allows you to set default parameters for virtual machines on a cluster basis, including VM HA, VM Cross-Cluster HA, VM CPU Model, Hyper-V Virtualization, Video Card Type on Boot, and vNIC Multi-queue Upgrading.
	Lifecycle Management	Allows you to manage the lifecycle of hosts, including adding, enabling, disabling, reconnecting, powering on, powering off, entering into maintenance mode, modifying, and deleting hosts.
	Monitoring	Allows you to view visualized monitoring data of hosts, including CPU, memory, disk, and NIC.
	Associated Resources	Allows you to manage associated virtual machines and Kernel adapters on a host.

Type	Features	Description
		After you deploy SAN storage on a host, you can manage LUNs on the host and pass through them to virtual machines.
		Allows you to manage physical NICs and aggregated ports on hosts. Supports viewing LLDP information. Supports monitoring aggregated ports status to help users quickly identify faulty NICs.
		Allows you to manage the physical GPU devices detected on a host and pass through them with other peripheral devices (such as GPU graphics cards and GPU sound cards) to virtual machines.
		Allows you to generate virtual GPU devices from physical GPU devices and attach these virtual GPU devices to virtual machines.
		Allows you to manage the USB devices detected on a host and pass through them to virtual machines.
		Allows you to manage PCI devices detected on a host, edit the PCI allowlist, and pass through these PCI devices to virtual machines.
	IPMI Management	Supports remote access to hosts via IPMI after users provide correct IPMIM information.
	Power Control	Allows you to power on or off a host after the host is managed by IPMI.
	Web Terminal	Allows you to enter the web terminal of a host and perform operations on the host.
	SSH Password Modification	Allows you to change the SSH password of a host . The new password takes effect after the host automatically reconnects.
	Encrypted Storage of Password	Supports storing host password in an encrypted form to protect user data privacy.
	Intel EPT Hardware Assist	Allows you to enable Intel EPT hardware assist for Intel CPUs to improve CPU performance.
	Custom Tag	Allows you to customize tags for hosts so that you can locate them quickly.

Type	Features	Description
	Custom List	Supports customized display of specific columns to achieve personalized and efficient information presentation.
	CSV File Export	Allows you to export host information as a CSV table, which helps in statistical analysis and problem diagnosis.
Virtual Machine	Lifecycle Management	Allows you to manage the lifecycle of virtual machines, including creating, powering on/off, rebooting, resuming, pausing, force stopping, shut down, modifying, exporting virtual machines.
	Monitoring	Allows you to view visualized monitoring data of virtual machines, including CPU, memory, disk, and NIC.
	Custom Display	Supports a host view for easy checking of physical associations of virtual machines.
		Supports a group view for easy management of virtual machines with same attributes or used in the same scenarios.
	VM Console	Allows you to access virtual machines through terminals without using remote tools.
		Supports three types of console mode: SPICE, VNC, and SPICE+VNC.
		Allows you to paste command tests into VNC console, manage VM power in VNC console. The VNC console supports automatically adapting to browser resolution.
		Allows you to set the number of screens for SPICE console.
		Allows you to set the console password. You can customize the password strength in the system parameters.
		Supports real-time preview of the virtual machine console display, with a default refresh rate of every 5 seconds.
	VM CPU	Allows you to specify the total CPU cores and cores per socket.

Type	Features	Description
		Allows you to specify whether to select the same CPU model as the host for a virtual machine. This configuration makes the virtual machine inherit some or all host CPU features that suit your business needs.
		Allows you to set CPU resource priority to enhance the ability of certain virtual machines to compete for resources when host loads are high.
		Allows you to assign the virtual CPUs of a virtual machine to specific host physical CPUs, which improves VM performance.
		Allows you to configure EmulatorPin for a virtual machine so that all other threads than virtual CPU threads and I/O threads of a virtual machine are assigned to physical CPUs of the host.
	VM Memory	Allows you to set memory resource priority to enhance the ability of certain virtual machines to compete for resources when host loads are high.
	VM Disk	Allows you to specify disk storage locations to meet different performance requirements for virtual machines.
		Supports RDM disks for direct LUN device access by virtual machines.
		Supports simulating different disk bus types.
		Allows you to set disk cache mode to control whether the host page cache is used when writing data to disks from virtual machines.
		Supports multiple virtual machines sharing the same disk under distributed storage or SAN storage .
	VM NIC	Allows you to enable/disable NICs while virtual machines are powered on/off.
		Allows you to attach/detach NICs while virtual machines are powered on/off.
		Allows you to set the NIC model to meet different business needs.

Type	Features	Description
		Allows you to set the number of NIC queues to choose whether to use multiple queues for sending and receiving network packets, enhancing network bandwidth performance.
		Allows you to set a QoS limit on the NIC while virtual machines are powered on/off.
		Allows you to specify the IP address and MAC address for virtual machines.
		Supports automatically sending and reading network configurations, such as IP address, through VMTools.
		Allows you to enable the NIC anti-spoofing mode to provide protection against IP/MAC spoofing and ARP spoofing.
	VM CD/DVD Drive	Supports loading ISO image files to boot virtual machines through ISO drive. Supports loading multiple drives for a single virtual machine to improve deployment efficiency.
	VM Peripheral Device	Allows you to attach/detach physical GPU devices and virtual GPU devices while virtual machines are powered on/off.
		Allows you to attach/detach USB devices while virtual machines are powered on/off.
		Supports four types of graphics card: vga, virtio, qxl, and cirrus.
		Supports three types of audio card: HDA(ICH6), HDA(ICH9), and AC97.
		Allows you to specify whether to enable hot plugging of PCI devices for a virtual machine.
	VM Snapshot	Allows you to take a snapshot at specified time points to record the state of the virtual machine, which allows rollback in case of breakdowns.
		Supports full revert (revert the VM data and disk order) and custom revert (revert only specified disks).

Type	Features	Description
		Supports VM auto boot after restoring from snapshots.
	VM Clone	Clones entire virtual machine data. Supports various cloning methods: Full Clone and Instant Full Clone.
		Allows you to specify storage allocation policy, including system allocation and manual allocation.
	VM Image	Allows you to create virtual machine images to facilitate customized batch creation of virtual machines.
	VM Template	Supports cloning a virtual machine into a template and converting a virtual machine to a template.
		Supports converting a template to a virtual machine and creating new virtual machines from a template.
		Supports modifying template configurations.
	VM Specification	Allows you to create VM customization specifications to avoid hostname or SID conflicts in VM batch deployment.
	VM Migration-Change Host	Allows you to migrate virtual machines to other hosts. Supports hot migration and cold migration.
		Hot migration: Copies CPU-related register states and memory. Supports local storage, NFS storage, distributed storage, and SAN storage.
		Cold migration: Supports local storage.
		Allows you to enable auto-converge to improve the success rate of the migration, if the migration is blocked because the virtual machine has been high-loaded for a long time.
	VM Migration-Change Data Storage	Allows you to migrate virtual machines to other data storage.
		Supports hot and cold migration across SAN storage and hot migration between SAN storage and distributed storage.
		Allows you to migrate the entire virtual machines (virtual machine and its data disks, except for shared disks).

Type	Features	Description
	VM Migration-Change Host and Data Storage	Allows you to specify storage pools for disks when hot migrating from SAN storage to distributed storage.
		Allows you to migrate virtual machines to other hosts and data storage. Supports hot migration and cold migration.
		Supports VM hot migration across the same type of data storage, including distributed storage - distributed storage, NFS storage - NFS storage, and SAN storage - SAN storage.
		Supports VM hot migration across different types of data storage, including distributed storage - SAN storage, local storage - SAN storage, local storage - distributed storage, local storage - NFS storage, SAN storage - NFS storage, and distributed storage - NFS storage.
		Supports VM cold migration across the same type of data storage, including distributed storage - distributed storage, and NFS storage - NFS storage.
		Allows you to enable auto-converge to improve the success rate of the migration, if the migration is blocked because the virtual machine has been high-loaded for a long time.
		Allows you to clean up the original data to release storage space, if data integrity is confirmed after storage migration.
	VMTools	Supports installing VMTools for virtual machines, including QEMU Guest Agent (QGA), Cloudbase-Init, advanced monitoring agent, and Virtio.
	VM HA	Supports automatic reboot of virtual machines when hosts fail. You can view the reboot progress on the UI.
		Allows you to globally control HA functionality through HA policies.
		Allows you to set a VM cross-cluster HA policy.

Type	Features	Description
	VM Security	Allows you to configure and manage standard key providers and native key providers.
		Supports key rotation for business or compliance requirements.
		Allows you to add a virtual Trusted Platform Module (vTPM) to provide enhanced security for the virtual machine operating system.
		Supports Secure Boot to enhance virtual machine security and infrastructure trust.
	Register VM	Allows you to register a virtual machine configuration file and data file to create a corresponding virtual machine record on the platform. After registration, the virtual machine is managed uniformly by the platform.
	Rest System	Allows you to reset a virtual machine to the initial state of the VM image. All data in the system disk will be overwritten.
	Change Owner	Allows you to change the owner of a virtual machine. After the modification, the new owner has full permissions over the virtual machine.
	Change VM Group	Allows you to change the group that virtual machines belong to for easier classification and management.
	Custom Tag	Allows you to customize tags for virtual machines so that you can locate them quickly.
	Set Hostname	Allows you to set the hostname when you create a virtual machine.
	VM Scheduling Group	Allows you to join in or exit from a VM scheduling group so as to associate with or disassociate from related VM scheduling policies. This way, you can manage the distribution of virtual machines on hosts.
	SSH Key Injection	Supports password-free login for Linux virtual machines by injecting an SSH public key.
	Import User Data	Allows you to import user data. You can upload user-defined parameters or scripts to customize

Type	Features	Description
		configurations for virtual machines or to accomplish specific tasks.
	Change VM Password	Allows you to change the password of a Linux or Windows running virtual machine.
	VM Boot Options	Allows you to set the boot order for virtual machines , supporting booting from disk, CD/DVD drive, or network.
		Supports two BIOS modes, including Legacy and UEFI.
		Allows you to set the BIOS post delay.
	Hide KVM Virtualization Flag	Allows you to disable the hypervisor for a virtual machine to make certain applications to skip their virtualization detection on this virtual machine.
	VMware I/O Port Simulation	Allows you to set whether to allow KVM virtual machines to emulate I/O ports in VMware environment, enabling KVM virtual machines to use the VMware I/O port standard for compatibility with VMware environments.
	Hyper-V	Allows you to enable Hyper-V for a Windows virtual machine.
	Export OVA Template	Supports exporting virtual machines as OVA templates.
	Custom List	Supports customized display of specific columns to achieve personalized and efficient information presentation.
	Export CSV File	Allows you to export virtual machine information as a CSV table, which helps in statistical analysis and problem diagnosis.
	Resource Deletion Protection	Displays warnings of the consequences on the UI and asks for confirmation before the deletion is completed.
		Provides three deletion policies to lower risks caused by misoperations. The policies include Direct, Delay (default), and Never.
Bare Metal Management	Bare Metal Template	Supports system template and custom template based on how the bare metal template is added.

Type	Features	Description
		Supports the following operating systems: custom platform OS and mainstream Linux distributions (RHEL/CentOS series, Debian/Ubuntu series, and SUSE/openSUSE series).
		Allows you to view the template content and download a bare metal template.
	Bare Metal Cluster	Allows you to manage the lifecycle of bare metal clusters, such as creating and deleting bare metal clusters.
		Allows you to attach/detach a deployment server to/from a bare metal cluster.
		Allows you to attach/detach distributed switches to/from a bare metal cluster.
	Deployment Server	Allows you to specify an independent server as the deployment server to provide PXE services and console proxies for bare metal chassis.
	Bare Metal Chassis	Supports two types of addition: manual addition and template import. You can add up to 500 bare metal chassis at a time.
		Allows you to manage the lifecycle of bare metal chassis, such as adding, enabling, disabling, powering on, powering off, rebooting, and deleting bare metal chassis.
		Allows you to automatically or manually obtain the hardware information of a bare metal chassis.
		Allows you to launch the console of a bare metal chassis and jump to its IPMI management page.
		Allows you to view the hardware configuration of a bare metal chassis.
	Bare Metal Instance	You can add up to 50 bare metal instances at a time.
		Allows you to select images (in ISO format and are not live CDs) to deploy operating systems for bare metal instances.
		Allows you to achieve unattended batch installation of bare metal instance operating systems with

Type	Features	Description
		preconfigured files generated from the bare metal template.
		Allows you to configure business networks for a bare metal instance.
		Allows you to manage the lifecycle of bare metal instances, such as creating, starting, stopping, rebooting, deleting, recovering, and expunging bare metal instances.
		Allows you to launch the console of a bare metal instance.
		Allows you to customize tags for bare metal instances so that you can locate them quickly.
		Supports visualized monitoring: displays the bare metal instance data such as CPU, memory, disk I/O , disk size, and NIC I/O.
		Allows you to centrally view the resources associated with a bare metal instance, such as NICs and disks.
Image Storage	Standalone Image Storage	Store image files through image slices and support incremental storage.
		Allows you to manage the lifecycle of standalone image storage, including adding, enabling, disabling , reconnecting, modifying, and deleting standalone image storage.
		Allows you to obtain the existing image files under the mount path of the standalone image storage.
		Allows you to specify a data network for a standalone image storage for data communication with hosts.
		Supports image synchronization between different standalone image storage on the same management node, and allows you to specify an image synchronization network for standalone image storage.
		Allows you to clean up invalid data stored in standalone image storage to release spaces.

Type	Features	Description
		Allows you to change the password for a standalone image storage.
		Allows you to centrally manage images in a standalone image storage.
		Monitors and displays the percentage of used capacity of a standalone image storage.
	Distributed Image Storage	Store image files through distributed block storage.
		Allows you to manage the lifecycle of distributed image storage, including adding, enabling, disabling , reconnecting, modifying, and deleting distributed image storage.
		Allows you to add multiple monitoring nodes and manage all the monitoring nodes centrally.
		Supports specifying image cache pools.
		Allows you to specify a data network for a distributed image storage for data communication with hosts.
		Allows you to centrally manage images in a distributed image storage.
		Allows you to clean up the original data preserved after migration across distributed image storage.
		Monitors and displays the percentage of used capacity of a distributed image storage.
	Image Management	Allows you to manage the lifecycle of images, including adding, modifying, deleting images.
		Supports two types of image: system image and disk image.
		Supports four types of image format: qcow2, iso, vmdk, and raw.
		Supports two types of image upload method: an URL or local browser.
		Supports MD5 verification to ensure that uploaded images are intact and undamaged.
		Supports exporting images.

Type	Features	Description
Data Storage		Allows you to migrate an image to another distributed image storage.
		Allows you to set the sharing mode of an image, including share globally, share to users/user group, and not share.
		Supports resource deletion protection and provides deletion policies.
	Local Storage	Allows you to add local storage by using free disks or local directory.
		Allows you to manage the lifecycle of local storage, including adding, enabling, disabling, reconnecting, entering maintenance mode, and deleting local storage.
		Allows you to manage associated virtual machines, disks, hosts, and clusters.
		Monitors and displays the storage utilization, storage distribution, and storage allocation ratio of local storage.
	SAN Storage	Uses shared block devices and supports setting up SAN storage through iSCSI storage, FC storage, or NVMe storage.
		Allows you to manage the lifecycle of SAN storage, including adding, enabling, disabling, reconnecting, entering maintenance mode, and deleting SAN storage.
		Allows you to manage the associated virtual machines, disks, clusters, shared LUNs.
		Supports two provisioning methods: thick provisioning and thin provisioning.
		Allows you to specify a storage network to check the virtual machine health status.
		Allows you to add multiple LUNs and refresh the storage capacity to view its changes when expanding or replacing a LUN device.

Type	Features	Description
		Allows you to forcibly clean up the data in a block device, such as the signature in the file system, RAID, and partition table.
		Allows you to clean up the original data preserved after migration across SAN storage.
		Monitors and displays the storage utilization, storage distribution, and storage allocation ratio of SAN storage.
	Distributed Storage	Uses distributed block storage and supports authentication with keys.
		Allows you to manage the lifecycle of distributed storage, including adding, enabling, disabling, reconnecting, entering maintenance mode, and deleting distributed storage.
		Allows you to manage the associated virtual machines, disks, clusters, storage pools, and monitoring nodes.
		Allows you to specify storage pools when adding a distributed storage.
		Allows you to specify a storage network to check the virtual machine health status.
		Allows you to clean up the original data preserved after migration across distributed storage.
		Monitors and displays the storage utilization, storage distribution, and storage allocation ratio of distributed storage.
	NFS Storage	Uses network file system and supports customize mount parameters.
		Allows you to manage the lifecycle of NFS storage , including adding, enabling, disabling, reconnecting, entering maintenance mode, and deleting NFS storage.
		Allows you to manage the associated virtual machines, disks, and clusters.
		Allows you to specify a storage network to check the virtual machine health status.

Type	Features	Description
		Allows you to clean up the original data preserved after migration across NFS storage.
		Monitors and displays the storage utilization, storage distribution, and storage allocation ratio of NFS storage.
Storage Service /		Allows you to deploy distributed storage through GUI.
		Allows you to seamlessly take over existing distributed storage.
Network Resource	Distributed Switch	Provides a virtual switch device for unified network management and monitoring of virtual machines within a cluster.
		Automatically generates default distributed switch for managing management network traffic on hosts.
		Allows you to create distributed port groups based on the default distributed switch to achieve reuse of management network and business network.
		Allows you to manage the lifecycle of distributed switch, including creating and deleting distributed switch.
		Supports VLAN (802.1Q) layer 2 isolation.
		Allows you to manage the uplinks associated with hosts on a distributed switch.
		Allows you to manage the associated distributed port groups and clusters.
	Distributed Port Group	Distributed port group is the logical grouping of distributed switch ports for port configuration.
		Allows you to manage the lifecycle of distributed port groups, including creating, modifying, and deleting distributed port groups.
		Supports adding IPv4 or IPv6 networks.
		Supports enabling or disabling IP address management.
		Supports enabling or disabling DHCP service.

Type	Features	Description
		Allows you to customize the MTU to limit the size of network transmission packets.
		Allows you to manage network segments, DNS, and other resources of distributed port groups.
		Supports visual monitoring and list statistics of IP usage for distributed port groups, helping to improve IP planning efficiency.
	Kernel Adapter	Automatically generates a default Kernel adapter for managing and configuring the physical network.
		Supported service type: Management and Storage.
Network Service	Security Group	Provides network security controls for virtual machine NICs.
Network Topology	/	Displays complete link relationships and other information centered around the distributed switch.
		Supports refreshing to display the latest network topology.
		Supports exporting the network topology as a PNG image.
		Supports hiding or displaying virtual machines in the network topology.
		Supports highlighting selected resources and displaying resource information in tooltips.
		Supports canvas operations such as fitting to canvas, zooming in, zooming out, and full-screen mode.
		Supports searching for resources by name or UUID within the current topology view.
Reliability	MN Monitoring	Allows you to set up a dual management node environment through GUI. If either node fails, service will automatically switch over within seconds to ensure continuous availability.
		Allows you to view the management service status , including whether the monitor IP is reachable, whether the peer management node is reachable , whether the virtual IP is reachable, and the database status.

Type	Features	Description
	Dynamic Resource Scheduling	Displays the DRS list for the cluster to easily view the platform's DRS policies.
	VM Scheduling Policy	Supports four types of scheduling policies: VM Exclusive from Each Other, VM Affinitive to Each Other, VMs Affinitive to Hosts, and VMs Exclusive from Hosts.
		Allows you to manage the lifecycle of VM scheduling policies, including creating, enabling, disabling, modifying, and deleting VM scheduling policies.
		Allows you to group a set of virtual machines for unified scheduling based on business requirements and manage the lifecycle of VM scheduling groups.
		Allows you to group a set of hosts for unified scheduling based on business requirements and manage the lifecycle of host scheduling groups.
	HA Policy	Supports controlling virtual machine high availability globally through HA policies.
		Allows you to configure VM failover strategies based on the combination of management network connectivity, storage network connectivity, and business NIC status.
		Allows you to modify host error detection policies and HA advanced settings.
		Supports viewing and filtering HA migration logs.
Data Protection	Snapshot Management	Supports centralized management of virtual machine and disk snapshots.
		The snapshot management interface is divided into two parts, allowing for linked display of virtual machines and their corresponding snapshots.
		Supports sorting the virtual machine list by snapshot count and snapshot capacity.
		Allows you to manage the lifecycle of snapshots, including creating, reverting, deleting snapshots.
	Backup Management	Supports viewing protected resource backup data in a tree-like folder structure.

Type	Features	Description
		Supports two data recovery methods: overwrite recovery and new virtual machine from backup.
		Allows you to manage the lifecycle of backup plans, including creating, enabling, disabling, and deleting backup plans.
		Supports creating backup plans for new virtual machines and platform databases.
		Supports two types of backup storage: local backup storage and remote backup storage.
		Supports three ways of backup storage addition: reuse image storage, reuse host, and dedicated backup storage.
		Allows you to manage the lifecycle of backup storage, including adding, enabling, disabling, reconnecting, and deleting backup storage.
Monitoring and Alarm	Resource Monitoring	Supports visual display of performance monitoring charts for resources, including clusters, hosts, virtual machines, image storage, data storage, and distributed port groups.
	Alarm	Provides a variety of alarm metrics to meet monitoring and alerting needs for multiple resource types and application scenarios.
		Supports two types of alarm: resource alarms and event alarms.
		Provides default alarms and allows you to customize alarms.
		Allows you to manage the lifecycle of alarms, including creating, enabling, disabling, and deleting alarms.
		Supports three levels of alarm: emergent, major, and info.
		Allows you to enable alarm recovery notification for resource alarms as needed. If enabled, when a resource monitored by a resource alarm recovers from the alarmed status, the system receives a notification.

Type	Features	Description
	Message Template	Allows you to centrally manage the endpoints and alarm records of an alarm.
		Sends messages to endpoints by using a text template. Supports customizing alarm message templates.
	Endpoints	Allows you to manage the lifecycle of message templates, including creating, modifying, and deleting message templates.
		Provides a system default endpoint and allows you to manage the lifecycle of the system endpoint, including enabling and disabling system endpoints.
		Supports custom endpoints, including email, DingTalk, Lark, WeCom, HTTP application, and Microsoft Teams.
		Allows you to manage the lifecycle of custom endpoints, including creating, enabling, disabling, and deleting endpoints.
		Allows you to add/remove alarms to/from an endpoint and centrally manage these alarms.
		Allows you to centrally manage messages received by an endpoint.
	Alarm Message	Supports intuitive viewing and unified management of platform alarm messages to improve operational efficiency.
		Displays alarm messages of different emergency levels in the last seven days on a bar chart.
		Displays alarm messages of different resources in the last seven days on a pie chart.
		Allows you to view up to 1,000 alarm messages in the message list.
		Supports multiple filter rules, including resource type, time period, alarm level, message type, and read/unread status.
		Allows you to set a silence period for alarm messages. During the silence period, no alarm

Type	Features	Description
		messages will be generated. You can process the alarm information when you are convenient.
		Allows you to cancel the silence period for alarm messages.
		Allows you to view the details about an alarm.
		Allows you to export the alarm messages as a CSV table, which helps in statistical analysis and problem diagnosis, and allows you to export the filtered alarm messages.
Task and Event	Operation Task	Displays the historic operations performed on the platform.
		Supports multiple filter rules, including time period, task result, and operator.
		Allows you to view the details about an operation task.
		Allows you to export operation task logs in CSV format.
	HA Task	Allows you to view and manage triggered HA migration tasks.
		Supports multiple filter rules, including time period and task result.
		Allows you to view the details about an HA task.
		Allows you to export HA task logs in CSV format.
	Scheduling Task	Allows you to view and manage scheduling execution history, results, and times.
		Allows you to view the details about a scheduling task.
		Allows you to export scheduling task logs in CSV format.
	Event	Monitors and records all activities in the platform, which effectively ensures the security of the cloud environment.
		Supports multiple filter rules, including time period and task result.

Type	Features	Description
		Allows you to view the details about an event.
Log Collection	Collect Log	Allows you to collect platform logs and logs of various nodes on the platform that are generated in the specified time range.
	Manage Log	Allows you to collect, cancel collection, download, and delete logs.
Tag Management	/	Allows you to customize tags for resources and quickly locate resources by tag type and tag name.
		Supports admin tags and user tags.
		Allows you to manage the lifecycle of tags, such as creating and deleting tags.
		Allows admins to attach/detach tags to/from all resources in the platform and users to attach/detach tags to/from resources of users.
		Allows you to centrally manage resources with a tag attached.
Custom Attribute	/	Allows you to create custom attributes and quickly filter target resources through attribute key and attribute value.
		Supports adding custom attributes to the following resources: Global, Virtual Machine, Host, Data Storage, Distributed Switch, Distributed Port Group, and Bare Metal Instance.
		Allows you to manage the lifecycle of custom attributes, such as creating and deleting custom attributes.
		Allows you to centrally view the associated resources of custom values.
Migration Service	Migration Service	Supports complete migration of virtual machine systems and data from other virtualization platforms to the current platform. Supported source platforms include VMware vSphere and other KVM-based cloud platforms. A step-by-step wizard guides you through the end-to-end migration process.

Type	Features	Description
	Migration Resource	Allows you to manage the lifecycle of source platforms, such as adding and deleting source platforms.
		Allows you to manage the lifecycle of data gateways, such as adding and deleting data gateways.
	Migration Task	Allows you to manage the lifecycle of migration tasks, such as creating, editing configurations, synchronizing, testing, and deleting migration tasks.
System Management	User Management	A user is created by the admin or synchronized from an SSO authentication system and is managed by the admin. Resources created under a user are managed by the user.
		Allows you to manage the lifecycle of local users, including creating and deleting users. Allows you to manage the lifecycle of SSO users, including synchronizing and deleting users.
		Allows you to add an SSO server to the platform so as to integrate the SSO system and enable password-free login of related accounts in the system.
		Supports three types of SSO server: OIDC, AD, and LDAP.
		Allows you to manage the lifecycle of the SSO server, such as adding and deleting the SSO server.
		Allows you to set two-factor authentication for user login, view the two-factor QR codes of the user, and download the two-factor QR codes.
		Allows you to set and manage resource quota for users, including compute resources, storage resources, and network resources.
		Allows you to centrally manage the associated or shared resources of a user.
		Allows you to change a regular user to an admin user.

Type	Features	Description
	Console Proxy	Allows you to set a console proxy to log in to a virtual machine.
		Allows you to reconnect a console 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.
		Allows you to manage the lifecycle of local AccessKeys, such as generating, enabling, disabling, and deleting AccessKeys.
	Security Settings	Allows you to manage platform login policies, virtual machine security policies, and host security policies to ensure security.
	Certificate Management	Allows you to configure and manage a SSL certificate, including third-party certificate and system self-signed certificate.
	Log Server	A log server is used to collect logs of the management node. You can add a log server to the platform and use the collected logs to locate errors and exceptions. This improves your O&M efficiency.
	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.
	System Parameter	Allows you to configure settings that take effect on the whole platform.
		Allows you to reset to default settings with one click.
		Supports quick search and directory navigation to help you quickly locate target items.
	Theme and Appearance	Allows you to customize the theme and appearance of the platform.
		Allows you to set the global appearance (theme), titles (browser/login interface/platform interface), and monitor (title and appearance/data monitoring method).
		Allows you to reset to default settings with one click.
	License Management	Provides multiple licensing options and you can purchase according to your needs.

Type	Features	Description
		Provides three licensing agreements: Basic Trial, Basic Paid, and Advanced Paid.
		Supports two licensing methods: USB key and request key.
		Allows you to view the current license status and licensing records.
		Provides license expiration reminders when your license is about to expire, expired, or license quota exceeds.
UI Highlights	Quick Navigation	Provides a quick navigation entry, which is convenient for users to quickly locate and enter the required features and services.
	Global Search	Provides one-stop global search, allowing you to search for features and resources.
	Right-click Operation	Supports right-click to call out operation panel when the resource tree is expanded.
	Keyboard Shortcuts	Provides global shortcuts and page shortcuts, support keyboard combination keys for quick access to menus, enhancing user experience and operational efficiency.
	Internalization	Supports changing the UI language, including Simplified Chinese and English.
Installation	One-click Installation	Allows you to complete installing and deploying the platform from scratch with one simple command.
		Supports multiple installation modes: Management Node, Compute Node, Expert Node.
Upgrade	Seamless Upgrade	Support seamless upgrades from lower versions to higher versions.
	Incremental Upgrade	Support incremental upgrades to significantly increase the speed of updates.
	Deployment Environment Upgrade	Support upgrading the deployment environment via expert mode.

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 **OK**.

What's next

- To enhance platform security and prevent unauthorized access, the platform provides login policy settings. For more information, see [Security Settings](#).
- The platform supports both HTTPS and HTTP protocols for UI access. For more information, see [Certificate Management](#).
- The platform supports both account login and unified identity authentication. For more information, see [Identity and Access Management](#).

6 License Service

This section mainly introduces the license agreement for NexaSphere and the basic usage methods.

- [Overview](#)
- [Use License](#)

6.1 Overview

NexaSphere License is used to authorize users to operate platform resources. You can understand the license comprehensively from the following dimensions:

- [Licensing Methods](#)
- [Licensing Agreements](#)

Licensing Methods

NexaSphere offers two licensing methods: USB Key licensing and Request Key licensing.

- USB Key licensing is implemented by preloading the license onto hardware. The licensing process is as follows:
 1. Obtain the USB Key hardware that has the license preloaded.
 2. In a dual management nodes environment, insert the USB Key into any management node to complete the licensing. In a single-node environment, insert it into that management node.
- Request Key licensing is implemented by uploading a software file. The licensing process is as follows:
 1. Obtain the compressed package file containing the relevant license.
 2. Upload the license file to any management node of NexaSphere to complete the licensing.
In a single-node environment, upload it to that management node.

Licensing Agreements

NexaSphere offers three licensing agreements: Basic Edition - Trial, Basic Edition - Paid, and Advanced Edition - Paid. Specific descriptions of the licensing agreements are as follows:

- Basic Edition - Trial: Allows you to add 1 host for free to trial basic computing virtualization, storage virtualization, and network virtualization functions, with a trial period of 30 days. Official

after-sales technical support services are not provided during the trial period. This edition can be used for deploying a testing environment.

- **Basic Edition - Paid:** Allows for paid licensing based on physical CPUs or the number of hosts, providing basic computing virtualization, storage virtualization, and network virtualization functions. Official after-sales technical support services are available during the after-sales service period. This edition can be used for deploying a production environment.
- **Advanced Edition - Paid:** Allows for paid licensing based on physical CPUs or the number of hosts, providing complete computing virtualization, storage virtualization, and network virtualization functions. Official after-sales technical support services are available during the after-sales service period. This edition can be used for deploying a production environment.



Note:

For details on the functional differences provided by different licensing agreements, please consult official sales personnel.

6.2 Use License

If you have purchased the required license, you can refer to the following content for using the license, including USB Key licensing and Request Key licensing:

- [USB Key Licensing](#)
- [Request Key Licensing](#)
- [License Status Descriptions](#)

6.2.1 USB Key Licensing

You can refer to the following steps to understand the complete usage process of USB Key license:

1. [Install License](#)
2. [Update License](#)
3. [View Licensing Records](#)

Install License

To install the USB Key license for the first time:

1. Obtain the USB Key hardware that has the license preloaded.
2. Insert the USB Key into the management node. In a dual management nodes environment, it can be inserted into either management node.

3. In the upper left corner of the NexaSphere navigation menu, click **System Management** > **License Management** > **License** to enter the **License** interface. Check the status of the USB Key to ensure it is **Normal**.
4. Click the **Synchronize** button on the left side of **Upload License**.
5. Check the license status to ensure it is **Valid**, which indicates that the installation is complete.

**Note:**

In environments using USB Key licensing, there will be approximately 2 seconds of information reading time each time you enter the **License Management** interface, after which the USB Key related information will appear.

Update License

To expand or add licenses:

1. Insert the USB Key into the management node.
2. In the upper left corner of the NexaSphere navigation menu, click **System Management** > **License Management** > **License** to enter the **License** interface. Check the status of the USB Key to ensure it is **Normal**.
3. Click **Copy** USB Key ID, use that ID to apply for the license, and obtain the license file.
4. Return to the **License Management** interface, click **Upload License**, and upload the license file.
5. Check the license information to ensure all details are correct and that the license status is **Valid**, indicating that the update is complete.

View Licensing Records

If you have completed the installation or update of the license, you can click the **Licensing Record** tab on this page to view detailed authorization history, including: upload time, management type, authorization method, authorized item, authorized quota, status, license issuance time, and the expiration time of the license.

6.2.2 Request Key Licensing

You can refer to the following steps to understand the complete usage process of Request Key license:

1. [Install License](#)
2. [Update License](#)

3. [View Licensing Records](#)

Install License

To install the Request Key license for the first time:

1. Obtain the compressed package file containing the relevant license.
2. Upload the License file to the management node where the NexaSphere virtual IP is located to complete the authorization. In a single-node deployment, upload it to that management node.
3. Check the license status to ensure it is **Valid**, indicating that the installation is complete.

Update License

To expand or add licenses:

1. In the upper left corner of the NexaSphere navigation menu, click **System Management** > **License Management** > **License** to enter the **License** interface. Click **Download**.
2. Use the request code to apply for and obtain the compressed package file containing the relevant authorization License.
3. Upload the License file to any management node of the NexaSphere to complete the authorization. In a single-node deployment, upload it to that management node.
4. Check the license status to ensure it is **Valid**, indicating that the installation is complete.

View Licensing Records

If you have completed the installation or update of the license, you can click the **Licensing Record** tab on this page to view detailed authorization history, including: upload time, management type, authorization method, authorized item, authorized quota, status, license issuance time, and the expiration time of the license.

6.2.3 License Status Descriptions

During the use of licenses, different license states may occur. You can learn about them through these chapters:

- [USB Key Status and Abnormality](#)
- [License Status and Alerts](#)

USB Key Status and Abnormality

The USB Key has the following statuses.

Status	Description
Normal (Ready)	The USB Key status is displayed as normal when there is only one USB Key inserted into the management node and it is functioning properly.
Removed (Missing)	After the USB Key is removed, the status of the USB Key is displayed as removed. At this point, the platform will check whether the current License has expired. - If it has not expired, a temporary License will be generated, and the buffer duration will vary based on the removal situation: - If it is the first removal, the buffer License duration is 30 days. - If it is not the first removal: If the time since the last removal is less than 60 days, it will inherit the expiration time of the previous buffer License; if it is 60 days or more, the buffer License will be updated to 60 days.
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, it will lead to an abnormal state. The abnormal status of the USB Key does not affect the 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 has multiple USB Keys inserted, it will lead to a USB Key failure. In a single management node environment, if that management node has multiple USB Keys inserted, it will also lead to a USB Key failure.  Note: When the USB Key status is failed and the buffer License has expired, the authorization cannot be used.

License Status and Alerts

A license has the following statuses.

Status	Description
Valid	An active platform license is added, and the corresponding position will indicate "Valid".
Expired	If the platform license has expired, the corresponding position will indicate "Expired" and the functionality will be unavailable.

If an exception occurs to your license, the platform will remind you accordingly. The details are listed as follows.

Abnormality	Details
License to Expire	When the remaining validity of the license is less than 15 days or the 7x24 hour after-sales service is within 30 days of expiration, a banner will appear after logging into the platform indicating that it is about to expire.
License Expired	When the platform license has expired, logging into the platform will automatically redirect to the License Management interface to remind that it has expired.  Note: If your existing services on the platform are still running normally, do not perform any operations (such as reconnecting hosts/image storage/data storage) to avoid affecting service operation.
Insufficient License Quota	When updating the license authorization, if the actual resource usage exceeds the authorized quota, most platform features will be restricted. The specific restriction policies are as follows: • You can log into the platform, and you can obtain, refresh, and view resource status. • Operations related to reducing quotas are supported, such as deleting hosts, deleting clusters, modifying host status, and deleting virtual machines. • Operations related to license management are supported, such as downloading request key, uploading licenses, and deleting licenses. • All other operations are restricted. Note: • License authorization includes three methods: CPU sockets, CPU cores, and hosts. Either method leading to actual resource usage exceeding the authorized quota will restrict platform usage. • Once the management node is shut down, it cannot be restarted. You can only log into the management node console and update the license via the command-line mode.

7 Initialize NexaSphere

The initialization wizard helps you quickly set up necessary resources for your environment.

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.

7.1 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.

- Select the image path. Options include URL and local file.
- Enter a name for the image.
- Select the image type. Options include system image and disk image.
- Select the image format. Options include qcow2, iso, vmdk, and raw.
- Click **OK**.

9. Create a distributed port group.

- Enter a name for the distributed port group.
- 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.
- 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.

7.2 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:
 - **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.
 - **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:
 - **Database Backup**: Select a 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.

8 Platform Resources Definition

This chapter primarily introduces the definitions of basic resources such as computing, storage, and networking on the NexaSphere virtualization platform, as well as the relationships between these resources:

- [Resource Definition](#)
- [Resource Relationship](#)

8.1 Resource Definition

The detailed resource definitions provide comprehensive explanations of all computing, storage, and networking resources on the platform. This prepares you theoretically for efficient resource usage in your business practices. The resources covered in this chapter are listed below:

- [Compute Resource|Data Center & Cluster & Host](#)
- [Compute Resource|Image & Image Storage](#)
- [Compute Resource|Virtual Machine & VM Group](#)
- [Storage Resource|Data Storage](#)
- [Network Resource|Distributed Switch & Distributed Port Group](#)

8.1.1 Compute Resource|Data Center & Cluster & Host

This section introduces the following three resources:

- [Data Center](#)
- [Cluster](#)
- [Host](#)

Data Center

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.

A data center include the following core resources:

- **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:** 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.
- **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.
- **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.

Through the use of data centers, you can divide the platform resources into multiple largest management domains, offering three main advantages:

- Achieve resource isolation at the data center level. Physical sub-resources such as data storage and hosts within different data centers can be isolated from each other, ensuring maximum stability and fault tolerance.
- Data centers can be flexibly deployed based on data center construction, reducing network latency and improving access speeds.
- Each data center's resource usage can be monitored independently.

Cluster

As mentioned above, a cluster is a collection of hosts (with no limit on the number) organized according to business logic under a data center. To provide complete services for virtual machines, it must be combined with data storage and distributed switches. Therefore, when planning a cluster, ensure that:

- All hosts within the cluster have the same operating system.
- All hosts within the cluster have identical network configurations.
- All hosts within the cluster can access the same data storage.
- The cluster must be equipped with data storage and distributed switches.

Clusters enable the management and isolation of physical computing resources. Below are some of the functional characteristics of clusters:

- NexaSphere supports clusters for both x86 and ARM architectures, managing hosts of x86 and ARM architecture respectively.

- It supports dynamic resource scheduling at the cluster level, monitoring CPU or memory load conditions of hosts within the cluster, and dynamically adjusting the services of virtual machines running on the hosts to balance the cluster load and improve platform stability.
- It supports overcommitment ratio configuration for CPU and memory resources at the cluster level, enabling over-provisioning to increase the utilization rate of cluster computing resources.
- It supports configuring dedicated networks for clusters based on business scenarios, such as VDI networks and VM migration networks.

Host

As mentioned above, a host is a physical server that provides CPU, memory, local storage, and network resources to virtual machines.

8.1.2 Compute Resource|Image & Image Storage

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.

Image: An image is a template file used by virtual machines or disks. Images are categorized into system images and disk images.

8.1.3 Compute Resource|Virtual Machine & VM Group

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.

Functional Components

Operating System: Installing an operating system on a virtual machine is similar to installing it on a physical host. NexaSphere supports installing operating systems via images. You can upload the image to the corresponding image storage, and after creating a new virtual machine using this image, install the operating system on it. NexaSphere supports all mainstream operating systems, including Windows and Linux.

Hardware Devices: Virtual hardware devices work in virtual machines similarly to how physical hardware works on physical hosts. The host Hypervisor provides virtual machines with virtual CPU, memory, disks, NICs, GPU, USB devices, and more. You can customize the configuration of virtual machine hardware devices and their features, such as:

- **CPU:** Supports customizing the number of vCPU cores; setting CPU modes to inherit some or all CPU characteristics from the host; binding vCPUs to specific pCPUs; or configuring CPU hot-plug for online modification of the number of CPU cores.
- **Memory:** Supports customizing virtual machine memory; setting higher memory resource priority for critical virtual machines; or configuring memory hot-plug for online memory modifications.
- **Disk:** Supports customizing disk capacity for virtual machines; setting cache modes to enhance IO performance; and configuring disk bandwidth/IOPS limits.
- **NIC:** Supports customizing NIC models; setting multiple queues for sending and receiving network packets to improve network PPS and bandwidth performance; and configuring inbound/outbound bandwidth limits for NICs.
- **Peripherals:** Supports loading external devices, such as physical GPUs, virtual GPUs, USB devices, etc.

VMTools: VMTools is a collection of drivers and utilities that enrich virtual machine functionality, improve performance, and provide advanced monitoring capabilities for virtual machines. It is recommended to install this tool after creating a new virtual machine.

Feature Characteristics

NexaSphere offers a wide range of virtual machine features, allowing you to configure different features according to various environments, for example:

- Hardware Feature Configuration: Refer to [Hardware Devices](#).
- High Availability Configuration:
 - It provides high availability configurations at three levels: virtual machine, cluster, and global . This is used to set up automatic restarts when a virtual machine shuts down abnormally, ensuring business high availability.
 - Supports configuring cross-cluster high availability for virtual machines, enabling automatic migration to clusters with sufficient resources.
- Security Configuration:
 - Supports setting remote login console mode and password for virtual machines, as well as login authentication configurations, ensuring only users with the appropriate permissions and passwords can log in.
 - Supports installing VMTools on virtual machines, which can automatically restart or shut down the virtual machine in case of a fault.

- Supports creating snapshots for virtual machines before important operations, preserving the data state at specific time points for quick rollback in case of faults.
- Supports a recycle bin policy for virtual machines: deleted virtual machines enter the recycle bin from the UI. If a deletion was accidental, you can restore the virtual machine; if it's no longer needed, you can delete it permanently.
- Supports operation logs and auditing for virtual machines, facilitating security analysis, intrusion detection, resource change tracking, and compliance audits.
- Migration and Scheduling Configuration:
 - Supports manually changing the host and/or data storage where a virtual machine resides, supporting both cold and hot migrations.
 - Supports assigning specific host groups to virtual machines through scheduling policies to ensure high performance and high availability of services.
- Group Management: Virtual machine grouping is a sub-resource of the data center. You can organize virtual machines across the entire data center into multiple levels of groups for easier viewing and management.

8.1.4 Storage Resource|Data 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. It makes full use of the host's local storage resources.
- Network Shared Storage: A storage system used for remote storage of virtual machines and their application data, accessible concurrently by hosts over a network. NexaSphere supports NFS, distributed storage, and SAN storage.
 - NFS: Utilizes a Network File System for storage, supporting custom mount parameters.
 - Distributed Storage: Employs a distributed block storage method, supporting key-based authentication.
 - SAN Storage: Uses shared block storage, supporting both thin provisioning and thick provisioning as storage allocation strategies.

SAN storage can be built through iSCSI Storage, FC Storage, and NVMe Storage:

- iSCSI Storage (iSCSI Storage): A storage technology that allows SCSI commands to be sent over IP networks. It provides block-level data access between servers and storage devices.
- FC Storage (FC Storage): Refers to Fibre Channel storage, which is a high-speed network technology primarily used for storage networking. It supports high-performance and reliable connections between servers and storage systems.
- NVMe Storage (NVMe Storage): Non-Volatile Memory Express storage, which is designed to accelerate read/write speeds using PCIe buses for faster performance compared to traditional SSDs.

8.1.5 Network Resource|Distributed Switch & Distributed Port Group

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. You can configure VLAN as needed.

8.2 Resource Relationship

There are two types of relationships between platform resources:

- Subordinate Relationships: Analogous to interpersonal relationships in human society, these include parent-child, sibling, grandparent-grandchild, and friend relationships.

Specific definitions are as follows:

- Parent-Child Relationship: Resource A is the parent or child of resource B. For example, clusters and hosts, hosts and virtual machines. In each pair, the latter runs within the former.
- Sibling Relationship: Resources A and B have the same parent and thus are siblings. For example, clusters and distributed switches, clusters and data storage; both pairs share a common parent, which is the data center.
- Grandparent-Grandchild Relationship: Resource A is the direct grandparent or grandchild of resource B. For example, the data center is the parent of clusters, clusters are the parents of hosts, and hosts are the parents of virtual machines. Therefore, the data center is the grandparent of hosts, and clusters are the grandparents of virtual machines.

- **Friend Relationship:** Resources A and B do not have any of the above three relationships but need to collaborate under certain circumstances. For example, data storage and image storage must work together to provide services for clusters.
- **Quantitative Relationships:** Similar to quantity limitation relationships in human society, these include 1:n (one-to-many), n:1 (many-to-one), and n:n (many-to-many).

Specific definitions are as follows:

- **1:n:** Indicates that resource A can create/add/load multiple instances of resource B. For example, a cluster can add multiple hosts, and a distributed switch can load multiple clusters.
- **n:1:** Indicates that multiple instances of resource A can create/add/load into one instance of resource B. For example, multiple hosts can be added to the same cluster, and multiple clusters can load the same distributed switch.
- **n:n:** Indicates that resource A can create/add/load multiple instances of resource B, and at the same time, multiple instances of resource A can also create/add/load into one instance of resource B. For example, an image storage can load multiple data centers, and a data center can also load multiple image storages.

8.2.1 Platform Resource Relationship

The following table illustrates the subordinate and quantitative relationships between various fundamental resources on the platform from the perspective of a data center:

Table 8-1: Resource Relationships from the Data Center Perspective

Resource A	Resource B	Subordinate Relationship	Quantitative Relationship
Data Center	Image Storage	Friend	n:n
	Cluster	Parent-Child	1:n
	Host	Grandparent-Grandchild	1:n
	VM Group	Parent-Child	1:n
	Virtual Machine	Grandparent-Grandchild	1:n
	Data Storage	Parent-Child	1:n
	Distributed Switch	Parent-Child	1:n

Resource A	Resource B	Subordinate Relationship	Quantitative Relationship
	Distributed Port Group	Grandparent-Grandchild	1:n

Table 8-2: Resource Relationships from the Cluster Perspective

Resource A	Resource B	Subordinate Relationship	Quantitative Relationship
Cluster	Host	Parent-Child	1:n
	VM Group	Sibling	/
	Virtual Machine	Grandparent-Grandchild	1:n
	Data Storage	Sibling	<i>n:n</i>
	Distributed Switch	Sibling	n:n

Table 8-3: Resource Relationships from the Host Perspective

Resource A	Resource B	Subordinate Relationship	Quantitative Relationship
Host	Virtual Machine	Parent-Child	1:n

Table 8-4: Resource Relationships from the VM Group Perspective

Resource A	Resource B	Subordinate Relationship	Quantitative Relationship
VM Group	Virtual Machine	Parent-Child	1:n

Table 8-5: Resource Relationships from the Data Storage Perspective

Resource A	Resource B	Subordinate Relationship	Quantitative Relationship
Data Storage	Image Storage	<i>Friend</i>	/
	Distributed Switch	Friend	/

Table 8-6: Resource Relationships from the Distributed Switch Perspective

Resource A	Resource B	Subordinate Relationship	Quantitative Relationship
Distributed Switch	Distributed Port Group	Parent-Child	1:n

8.2.2 Cluster|Data Storage|Image Storage Relationship

Cluster and Data Storage

Clusters and data storage are both sub-resources under a data center, and they share a sibling relationship. Data storage must be attached to a cluster before it can provide storage services to the virtual machines within the cluster. The number of data storage devices that can be attached to a cluster varies depending on the storage type.

Resource A	Resource B	Quantitative Relationship
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)

Image Storage and Data Storage

Once an image storage is attached to the data center, it will provide image storage services to the clusters within the data center. Different types of image storage must be paired with corresponding

types of data storage to support operations such as creating new virtual machines within the cluster. The following table outlines the specific combination restrictions.

Image Storage Type	Data Storage Type	Support New 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
	nSDS Distributed Storage	Yes
	SAN Storage	No

9 Compute Management

This chapter mainly introduces how to use compute virtualization resources, including data centers, clusters, hosts, image storages, images, and virtual machines.

- [Data Center](#)
- [Cluster](#)
- [Host](#)
- [Image Storage](#)
- [Virtual Machine](#)

9.1 Data Center

The data center is the largest resource namespace within the platform. This section describes how to use data centers in the following chapters:

- [Data Center Basic Operations](#)

9.1.1 Data Center Basic Operations

You can understand the basic operations supported by the data center from the perspective of CRUD (Create, Read, Update, Delete).

- [Create a Data Center](#)
- [Edit a Data Center](#)
- [View a Data Center](#)
- [Delete a Data Center](#)

Create a Data Center

Navigate to the root node and click **New Data Center**. Set the following parameters to create a data center:

- **Name:** Name of the data center
- **Description:** Description of the data center

After clicking **OK**, the data center will be created.

Edit a Data Center

If you need to modify the name or description of an existing data center, you can do so on the corresponding data center page by clicking **More Actions****Edit Name and Description**. In the dialog box that appears, you can modify the relevant information.

View a Data Center

NexaSphere provides a data center resource topology diagram, which helps you quickly grasp summary information about all resources within a data center. You can view this information on the **Overview** page of the target data center.

Delete a Data Center

If you need to delete an existing data center, you can do so on the target data center page by clicking **More Actions****Delete**. This will delete the data center.

**Note:**

Deleting a data center will also delete clusters, hosts, networks, data stores, and all subordinate resources (such as virtual machines) within them. Please proceed with caution.

9.2 Cluster

A cluster is a sub-resource of a data center, consisting of one or more hosts. This section describes how to use clusters in the following two chapters:

- [Cluster Basic Operations](#)
- [Cluster DRS](#)

9.2.1 Cluster Basic Operations

You can understand the basic operations supported by clusters from the perspective of CRUD (Create, Read, Update, Delete).

- [Create a Cluster](#)
- [Edit a Cluster](#)
- [View Cluster Capacity Information](#)
- [Delete a Cluster](#)

Create a Cluster

The platform provides multiple entry points for creating clusters. You can create a cluster from the following two main entry points:

- In the navigation bar on the left side of the platform page, right-click the target data center and click **New Cluster**.
- In the navigation bar on the left side of the platform page, select the target data center. Then, on the right side of the platform page, click **Actions > New Cluster**, or on the **Cluster and Host** sub-page, click **New Cluster**.

NexaSphere supports the following major categories of information configuration:

Basic Information Configuration: Includes name and description, associated data center, CPU architecture, and dynamic resource scheduling configuration

- **Name:** Name of the cluster
- **Description:** Description of the cluster
- **Data Center:** Data center where the cluster is located
- **CPU Architecture:** Supports creating x86 or aarch architecture clusters. The architecture of hosts added to the cluster must match that of the cluster.
- **DRS:** Monitors CPU/memory load of all hosts in the cluster and dynamically adjusts the virtual machine services running on hosts based on scheduling policies. This feature is disabled by default. If enabled, manual scheduling and automatic scheduling are supported. For details, refer to [Cluster DRS](#)

Advanced Configuration: Includes cluster networking, cluster over-provisioning, host settings within the cluster, and virtual machine settings within the cluster:

- **Cluster Network:** Configure dedicated VDI network CIDR and VM migration network CIDR for the cluster. If not set or ineffective, the management network is used by default.
- **Cluster Overcommitment:** Controls the number of virtual CPUs and the amount of virtual memory allocated to virtual machines:
 - Allocatable Virtual CPUs = Physical CPU Total Threads in the Cluster × CPU Overcommit Ratio
 - Allocatable Virtual Memory Capacity = (Physical Memory Capacity – Reserved Capacity) × Memory Overcommit Ratio
- **Host Settings:** Customize host characteristics within the cluster:
 - **Host CPU Model Check:** Whether the system checks for consistency between the CPU model of the source host and the CPU models of hosts in the current cluster during live migration or host addition. Default is disabled. If enabled, inconsistent models prevent the corresponding operation.

- **ignore_msrs Option:** Whether the ignore_msrs option is enabled in the host KVM kernel module. Default is disabled.
- **Host Zero Copy:** Whether the host CPU enables Zero Copy. Default is not enabled. Enabling reduces the number of copies of data between kernel space and user space, reducing CPU usage time and improving virtual NIC performance.
- **Huge Pages:** Whether hosts in the cluster enable huge pages (each page is 2MB) and allocate huge page space to virtual machines. Default is not enabled.
- **Host Reserved Memory:** The amount of memory reserved on the KVM host to ensure the host system can run normally. When huge pages are enabled in the cluster, hosts must reserve at least 4GB of memory.
- **Virtual Machine Settings:** Customize virtual machine characteristics within the cluster:
 - **VM HA:** Automatic restart mechanism for virtual machines after shutdown. Default is enabled. Details can be found in [VM HA](#).
 - **VM Cross-Cluster HA:** In shared storage (NFS, SAN, and distributed storage) scenarios, whether high availability migration across clusters is supported when changing hosts or when hosts enter maintenance mode.
 - **VM CPU Model:** Whether the virtual machine CPU model is set to match the host CPU model, inheriting some or all of the host's CPU features to meet different business needs. Default is unset (none) for x86_64 architecture, and passthrough for aarch architecture.
 - **Hyper-V Virtualization:** Whether the virtual machine enables Hyper-V emulation functionality. Default is disabled. Mainly used for nested virtualization scenarios with Windows systems.
 - **Video Card Type on Boot:** The default graphics card type used when starting virtual machines, providing basic/high-definition/high-performance video functionality. Default is vga for x86 architecture; virtio is the default and only option for aarch architecture.
 - **vNIC Multi-queue Upgrading:** Uses multiple queues for virtual machine NICs to receive and transmit network packets to improve network PPS and bandwidth performance.

After clicking **OK**, the cluster will be created.

Edit a Cluster

If you need to modify the name or description of an existing cluster, you can do so on the target cluster page by clicking **ActionsEdit Name and Description**. In the dialog box that appears, you can modify the corresponding information.

If you need to modify advanced settings of an existing cluster, such as cluster networking settings, over-provisioning settings, host settings within the cluster, or virtual machine settings, you can do so on the target cluster page by clicking **ActionsModify Configuration**, or on the **Advanced Settings** sub-page.

View Cluster Capacity Information

If you need to check the usage and allocation of physical CPU and memory resources for the cluster, as well as the usage and allocation of all sub-resources under the cluster, go to the **Overview** details page. For more information, please refer to [Capacity Monitoring](#).

Delete a Cluster

If you need to delete an existing cluster, you can do so on the target cluster page by clicking **ActionsDelete**. You can also delete clusters in bulk on the **Cluster and Host** page within the data center resources.



Note:

Deleting a cluster deletes all hosts within the cluster. If local storage is loaded, it will also delete all virtual machines and snapshots on the hosts. Proceed with caution.

9.2.2 Cluster DRS

Dynamic Resource Scheduling (DRS): Monitors CPU or memory load on hosts at the cluster level and dynamically adjusts virtual machine services running on hosts according to configured scheduling policies.

NexaSphere supports both manual and automatic scheduling strategies, both of which can balance cluster loads and effectively improve platform stability:

- Manual scheduling strategy provides scheduling suggestions, allowing you to manually migrate virtual machines according to the suggestions.
- Automatic scheduling strategy, where the system automatically performs resource scheduling based on scheduling algorithms.

The basic process for using dynamic resource scheduling is:

1. Configure the dynamic resource scheduling policy.
2. Execute dynamic resource scheduling operations.

Configure DRS

Prerequisites: To enable or use the dynamic resource scheduling feature, ensure that the cluster meets the following conditions:

- The cluster only attaches nSDS distributed storage or SAN storage.
- All hosts in the cluster have consistent CPU models.

Configure Dynamic Resource Scheduling:

The platform provides multiple entry points for configuring dynamic resource scheduling policies. You can configure dynamic resource scheduling policies from the following two main entry points:

- In the target cluster **DRS** tab, click **Enable** to configure the settings.
- In the navigation bar on the left side of the platform page, click **Business Reliability > DRS Policy**, then in the target cluster, click **Actions > Modify Policy** to configure the settings.

You can configure dynamic resource scheduling policies as follows:

- **Resource Type:** Default is compute resources, currently not modifiable.
- **Scheduling Policy:** Supports manual and automatic scheduling modes:
 - **Manual Scheduling:** After the CPU utilization or memory utilization of hosts in the cluster reaches the specified threshold, you manually execute resource scheduling based on the scheduling suggestions.
 - **Automatic Scheduling:** After the CPU utilization or memory utilization of hosts in the cluster reaches the specified threshold, the system automatically executes resource scheduling based on scheduling algorithms.
- **Monitoring Items:** Select the monitoring items for hosts, including: CPU utilization, memory utilization, CPU or memory utilization
 - **CPU Utilization:** Define the trigger conditions for the CPU utilization monitoring item.
 - **Memory Utilization:** Define the trigger conditions for the memory utilization monitoring item.
- **Duration:** Define the duration for the threshold, units include: seconds, minutes, hours.
- **DRS VM Migration Concurrency:** Number of virtual machines concurrently migrated during dynamic resource scheduling. Default is 1.
- **DRS Cluster Scan Interval:** Time interval for scanning the cluster balance status during dynamic resource scheduling. Default is 10 minutes.



Note:

If any host in the cluster reaches the threshold for the monitoring item trigger condition and meets the duration, it can be determined that the cluster state is imbalanced, and scheduling suggestions will be given.

Execute Dynamic Resource Scheduling Operations

After configuring dynamic resource scheduling, you can perform the following operations:

- Manually scan the cluster balance status.
- If set to manual scheduling, you can migrate virtual machines to recommended hosts based on scheduling suggestions to balance the cluster load.
- You can view scheduling execution history, results, and times in **O&M Management > Tasks > Scheduling Tasks**. By default, the last 7 days of data are displayed. You can customize the time period to view execution history and search execution history by virtual machine UUID.
- You can reconfigure the dynamic resource scheduling policy.
- You can disable cluster dynamic resource scheduling.



Note:

This may result in an inability to balance host loads in high-load scenarios, affecting business performance. Proceed with caution.

9.3 Host

A host is a sub-resource of a cluster and can run one or more virtual machines. This section describes how to use hosts in the following chapters:

- [Host Basic Operations](#)
- [Host Hardware Device](#)

9.3.1 Host Basic Operations

You can understand the basic operations supported by hosts from the perspective of adding, deleting, modifying, and querying.

- [Add a Host](#)
- [Modify a Host](#)
- [Access a Host](#)
- [View a Host](#)
- [Delete a Host](#)

Add a Host

The platform provides multiple entry points to add hosts. You can add one or multiple hosts from the following two main entry points:

- In the navigation bar on the left side of the platform page, right-click the target cluster and click **Add Host**.
- In the navigation bar on the left side of the platform page, select the target cluster. Then, on the right side of the platform page, click **Actions > Add Host**, or in the **Hosts** sub-page, click **Add Host**.

NexaSphere supports configuration of the following three major categories of information:

Basic Information: includes name description, associated cluster, and labels

- **Name:** host name
- **Description:** host description
- **Cluster:** cluster where the host resides
- **Tag:** supports binding one or more labels to identify different hosts. For more details, see [Tag Management](#)

Host Information: includes addition method, IP address range, and SSH configuration:

- **Addition Method:** add a single host or multiple hosts. When adding multiple hosts, ensure all hosts have the same SSH configuration.
- **IP Address:** enter the IP address or IP address range of the host based on the addition method.
- **SSH Port:** SSH port of the host, default is 22.
- **SSH Username:** username for the host. Default is root.
- **SSH Password:** password for the SSH username.

Other Information Configuration: includes IOMMU enablement status and Intel EPT hardware assistance

- **Scan Host IOMMU Setting:** whether the IOMMU (Input/Output Memory Management Unit) function is enabled, used for passthrough of external devices and virtualization scenarios. By default, it is disabled on x86 architecture and enabled on ARM architecture where the BIOS IOMMU is called SMMU.



Note:

Before enabling, make sure the IOMMU option is enabled in the host's BIOS.

- **Intel EPT Hardware Assist:** whether Intel EPT hardware assistance is enabled on Intel CPUs to improve CPU performance. Default is enabled.

After clicking **OK**, the creation process is complete.

After the first host is added to the cluster, NexaSphere automatically creates a default distributed switch, default distributed port group, and default Kernel adapter based on the related configurations of this host for centralized management of the host's management network. For more information, see [Network Resource](#).

Modify a Host

If you need to modify the name or description of an existing host, on the target host page, click **ActionsEdit Name and Description**, and modify the corresponding information in the pop-up window.

If you need to modify the IOMMU enablement status and Intel EPT hardware assistance settings of an existing host, on the target host page, click **ActionsModify Configuration**, and modify the corresponding information in the pop-up window.

Access a Host

You can access the host system through the following three methods:

- **Webshell Terminal Access:** you can directly access the host system by clicking on the small terminal window of the target host, or by clicking **Actions > Enter Web Terminal**.
- **SSH Access:** you can log in using remote login software by entering the SSH information entered when adding the host. To modify SSH information, click **Actions > Update SSH Information**.
- **IPMI Access:** if you have managed the host through IPMI, you can access the host using IPMI management software. To modify IPMI information, click **Actions > Update IPMI Information**.

View a Host

If you need to check the usage and allocation of CPU and memory resources for the host, as well as the usage and allocation of all virtual machines under the host, go to the host's **Overview** details page. For more information, see [Capacity Monitoring](#).

If you need to check the usage trends of CPU, memory, disk, and NIC resources for the host and its virtual machines over time, go to the host's **Monitoring** tab. For more information, see [Resource Performance Monitoring](#).

Delete a Host

If you need to delete an existing host, on the target host page, click **ActionsDelete**, and the host will be deleted. You can also delete hosts in batches on the data center resources **Cluster and Host** page or the cluster resources **Host** page.



Note:

- If the cluster to which the host belongs has loaded shared storage, this operation will stop all VMs on the host. VMs with high availability enabled will automatically migrate to other hosts within the cluster with sufficient resources and restart.
- If the cluster to which the host belongs has loaded local storage, this operation will delete all VMs and disks on the host. Proceed with caution.

9.3.2 Host Hardware Device

After adding the host to the NexaSphere platform, you can enter the host's **Hardware Device** page to view and manage the host's hardware and devices:

- [Host NUMA Topology](#)
- [Host NIC](#)
- [GPU Device](#)
- [USB Device](#)
- [PCIe Device](#)

9.3.2.1 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

- Non-Uniform Memory Access (NUMA): 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 primarily consists of one or more physical CPU cores (pCPU) 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 primarily consists of one or more virtual CPU cores (vCPU) 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.

Functionality Principle

After adding the host, NexaSphere supports viewing the host's pNUMA topology and configuring vNUMA for virtual machines running on the host based on this topology.

NexaSphere Virtual machine vNUMA configuration is achieved through CPU pinning, which strictly associates the virtual machine's vCPU with the host's pCPU, allocating specific pCPUs to the virtual machine. During vNUMA configuration, all vCPUs of the virtual machine are pinned to pCPUs, and each vCPU's pinned pCPUs are located within the same pNUMA node.

After vNUMA configuration, the virtual machine directly passes through the associated host pNUMA node topology, generating one or more vNUMA nodes that form the virtual machine's vNUMA topology. Virtual machine vCPUs prioritize accessing local memory within the same node based on the vNUMA topology.

pNUMA Topology

Go to the target host's **Overview**, and click **View pNUMA Topology** in the hardware overview information box. NexaSphere host pNUMA topology information is as follows:

- Displays all pNUMA nodes of the host and the virtual machine information associated with each node.
- Total memory is the local memory of the pNUMA node that can be directly accessed by the pCPU.
- Free memory is the local free memory of the pNUMA node that can be directly accessed by the pCPU.
- Both total memory and free memory are based on the actual hardware physical memory capacity of the pNUMA node.

9.3.2.2 LUN

On the host **Hardware Device** tab, click **LUN** to check the block device on the host.

- [View SCSI Devices](#)
- [View NVMe Devices](#)

View SCSI Devices

You can view basic information and path information about SCSI block devices scanned on the host.

Procedure:

1. Navigate to **Inventory > Host and VM**.
2. Select the target host, then click the host name to enter the details page.
3. Click **Hardware Device > LUN > SCSI Device**.

On the **SCSI Device** tab, view the supplier, model, capacity, WWN, WWID, number of mounted virtual machines, type, and source information for the block devices.

4. Click the name of the target block device, then click **Paths** to view all available paths for the block device and the status of each path.

View NVMe Devices

You can view NVMe LUNs connected to the host via network protocols. Local NVMe PCIe disks on the host can be viewed in the **Physical Disk** list.

Procedure:

1. Navigate to **Inventory > Host and VM**.
2. Select the target host, then click the host name to enter the details page.
3. Click **Hardware Device > LUN > NVMe Devices**.

On the **NVMe Devices** tab, view the supplier, model, WWN, capacity, WWID, and type information for the block devices.

9.3.2.3 Host NIC

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

- [Physical NIC - Standard Configuration Changes](#)
- [Physical NIC - SR-IOV Virtualization](#)

- [Physical NIC - View and Maintain LLDP Information](#)
- [Bond - View and Configure](#)

Physical NIC - Standard Configuration Changes

You can perform editing or IP address modification operations on physical NICs of the host:

- **Edit:** Modify the description of the physical NIC.
- **Modify IP Address:** If the NIC has not been added to a bond interface and is not part of a distributed switch, you can modify the IP address and subnet mask of the NIC as needed.

Physical NIC - SR-IOV Virtualization

You can virtualize and partition a single physical NIC into multiple VF (Virtual Function) NICs based on the SR-IOV specification, which can then be directly assigned to virtual machines. This allows for near-native I/O performance and reduces consumption of host CPU resources.

Prerequisites

- Ensure that the physical NIC supports SR-IOV partitioning.
- Ensure that the BIOS of the host containing this physical NIC has Intel VT-d/AMD IOMMU and SR-IOV features enabled.
- Ensure that the IOMMU readiness status of the host containing this physical NIC is **Available**.

Procedure

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

Notes

- If the physical NIC is already configured in a bond, continuing to use the SR-IOV feature may affect communication between VF cards and vNIC cards. It is recommended to first perform SR-IOV partitioning on the physical NIC, then configure the bond from the distributed switch, selecting only one partitioned physical NIC per host.
- If VF cards are currently being used by virtual machines, turning off the **SR-IOV Status** switch will simultaneously unload the related NICs from the virtual machines.

- Virtual machines that are powered on and have loaded VF cards do not support migration operations. You need to power off the virtual machine or first unload the VF cards before migration can occur.

Physical NIC - View and Maintain LLDP Information

View Peer Device Information

Using the Link Layer Discovery Protocol (LLDP), you can identify the physical switch port of the distributed switch to which the physical NIC is connected. Before viewing the peer device information, ensure that your NIC supports LLDP and that the peer switch device has LLDP enabled.

- 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 or x86_KylinV10P3
- ARM: arm_KylinV10P3 or H22e
- Procedure:
 1. Click on the physical NIC name to open the details page.
 2. Click **LLDP** to view the peer device information. For the interpretation of LLDP TLV units, refer to [Table 9-1: Appendix: LLDP Field Definitions](#).

Modify LLDP Mode

Using the Link Layer Discovery Protocol (LLDP), you can obtain peer device information connected to the port or send your own device information to neighboring devices directly connected to you, allowing for link communication status queries and judgments.

Steps:

1. Click on the physical NIC name to open the details page.
2. In the NIC details page, click **LLDP**.
3. On the **LLDP** tab, click **Modify**.
4. In the **Modify NIC LLDP Mode** dialog, select the LLDP mode from the dropdown options.
 - **Receive Only**: The default mode, which only parses and displays the peer LLDP information received on this port.

- **Send Only:** Sends the LLDP information from this port but does not parse the received LLDP information. Peer device information cannot be viewed in this mode.
- **Send and Receive:** Parses and displays the peer LLDP information received on this port while also sending the LLDP information from this port to the connected peer device.
- **Disabled:** Does not parse received LLDP information and does not send any LLDP information from this port. Peer device information cannot be viewed in this mode.

5. Click **OK**.

Table 9-1: Appendix: LLDP Field Definitions

Field	Description
Device ID	Chassis ID, which is the bridge MAC address of the sending device.
Port ID	Port ID, which identifies the port.
Management Address	Management Address, which is the management address of the sending port.
TTL	Time to Live, indicating how long this device's information remains on neighbor devices.
Port Description	Port Description, which provides details about the port.
System Name	System Name, which identifies the name of the device.
System Description	System Description, providing information about the system.
System Capabilities	System Capabilities, indicating the primary functions of the system and those that are in use.
VLAN ID	Primary VLAN ID of the port.
Aggregation Status	Link Aggregation, indicating whether the port supports link aggregation and if it has been enabled.
MTU	Maximum Frame Size, which is the maximum frame length supported by the port, taken from the configured Maximum Transmission Unit (MTU) value.

Bond - View and Configure

If you associate a physical network port with a distributed switch for aggregation when creating the distributed switch, you can view the aggregation port on the corresponding host resource's physical NIC **Bond** page. Click the Refresh button to view the latest information, including the port's aggregation mode, aggregation port status, rate, associated distributed switch, IPv4

address, and creation time. If you need to manage the aggregation port configuration for this host, refer to [Manage Joined Hosts](#).

9.3.2.4 GPU Device

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

- [Physical GPU Device](#)
- [vGPU Devices](#)

Physical GPU Device

You can perform different operations on physical GPU devices based on various scenarios:

- **Enable and Disable Scenarios:** Enable & Disable
 - If you want to pass through the physical GPU device directly to a virtual machine, ensure that the device is enabled. Click the **Enable** button to enable it.
 - If you no longer wish to pass through the physical GPU device to any virtual machines, click the **Disable** button to disable it.



Note:

After disabling, the physical GPU device currently in use by a virtual machine will continue to function normally until it is uninstalled.

- **Virtualization Scenarios:** Virtualization Partitioning & Restoration
 - **Virtualization Partitioning:** Partition an unpassed-through physical GPU device into vGPU devices of specified specifications. Before partitioning, ensure the following conditions are met:
 - The physical GPU model supports virtualization partitioning.
 - The physical GPU is not passed through to any virtual machine.
 - The host BIOS has Intel VT-d / AMD IOMMU enabled, and the host kernel has IOMMU support enabled.
 - The host's IOMMU readiness status in the platform is set to **Available**.

Different manufacturers have slightly different methods for partitioning physical GPUs virtually:

- **NVIDIA:** Supports partitioning NVIDIA physical GPUs according to selected specifications individually.

- AMD: Supports partitioning all AMD physical GPUs on the current host simultaneously based on the selected quantity.
- **Virtualization Restoration:** Restore vGPU devices back to physical GPU devices. Before restoration, ensure that all vGPUs created from this physical GPU have been uninstalled from virtual machines. The method of restoring physical GPUs varies by manufacturer:
 - NVIDIA: Ensure that all vGPUs related to this NVIDIA physical GPU have been uninstalled from virtual machines before restoration.
 - AMD: Ensure that all AMD vGPUs on the current host have been uninstalled from virtual machines before restoring AMD vGPUs.

vGPU Devices

You can perform different operations on vGPU devices based on various scenarios:

- **Enable and Disable vGPU Device Scenarios:** Enable & Disable
 - If you want to pass through the vGPU device directly to a virtual machine, ensure that the device is enabled. Click the **Enable** button to enable it.
 - If you no longer wish to pass through the vGPU device to any other virtual machines, click the **Disable** button to disable it. After disabling, the vGPU device currently in use by a virtual machine will continue to function normally until it is uninstalled.

9.3.2.5 USB Device

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

You can perform different operations on USB devices based on various scenarios:

- **Device Renaming Scenario:** If you wish to rename a USB device to better align with your business needs, click the **Edit Device Name** button to rename it.
- **Enable and Disable Scenarios:**
 - If you want to pass through a USB device to a virtual machine, ensure that the device is enabled. Click the **Enable** button to enable it.
 - If you no longer want to pass through the USB device to any other virtual machines, click the **Disable** button to disable it. After disabling, the USB device currently in use by a virtual machine will continue to function normally until it is uninstalled.
- **VM Attach/Detach Scenarios:** Attach VM & Detach VM

- **Attach VM:** Pass the USB device directly to a virtual machine, supporting both Direct Connection and Forwarding modes.
 - Direct Connection: Attach the USB device from the host where the VM resides to the VM. This USB device must be unloaded if the VM is migrated.
 - Forwarding: Attach the USB device from any host within the data center where the VM resides to the VM. This USB device does not need to be unloaded if the VM is migrated.

When loading a USB device to a virtual machine, note the following:

- The same USB device can only be passed through to one virtual machine at a time.
 - A single virtual machine supports up to 1 USB 1.0 device, up to 6 USB 2.0 devices, and up to 4 USB 3.0 devices.
 - A running VM or a VM with local storage in a stopped state only supports loading available USB devices from the host where the VM resides; cross-host USB device loading is not supported.
 - A VM stopped on shared storage supports loading multiple USB devices from any host within its cluster.
- **Detach VM:** Detach the USB device from the virtual machine.



Note:

This action will interrupt read/write operations for the USB device, proceed with caution.

9.3.2.6 PCIe Device

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

You can enable passthrough functionality for PCIe devices on the host. Passthrough devices provide an effective way to use resources and improve environment performance.

- **Passthrough PCI Devices:** Displays PCIe devices with passthrough functionality enabled, available for virtual machine use.
- **All PCI Devices:** Displays all PCIe devices detected on the host. You can switch devices with a Configurable passthrough state to passthrough devices.



Note:

- Before switching a PCIe device to passthrough, ensure that IOMMU is enabled on the host and that the IOMMU readiness state is available.
- If a PCIe device is already loaded by a virtual machine, it cannot be switched to passthrough. Please unload the virtual machine before retrying.

9.4 Image Storage

Image storage acts as a repository for storing image templates containing operating systems, as well as disk images. Through image storage, you can share image files across multiple data centers.

This section introduces how to use image storage from the following chapters:

- [Image Storage Basic Operations](#)
- [Image Basic Operations](#)
- [Image Cross-Platform Usage](#)

9.4.1 Image Storage Basic Operations

9.4.1.1 Add a Standalone Image Storage

A standalone image storage stores image files through 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, set the following parameters to complete basic configurations:
 - a) Set basic information.
 - **Name**: Image storage name.
 - **Description**: Image storage description.
 - **Type**: Displays standalone image storage.
 - **Data Center**: Current data center.
 - b) Set image storage configurations.
 - **Image Storage IP**: IP address of the image storage server.

- **SSH Port:** Default port is 22. Supports setting SSH port.
- **Username:** Default is root user. Supports setting regular user.
- **Password:** Password for the user.

6. Click **Next**.

The system will test connectivity to the image storage IP.

7. Set the following parameters to complete storage configuration:

a) Set disk configurations.

- **Addition Method:** Supports **Free Disk** and **Local Directory** addition methods.

When selecting **Free Disk**, configure the following parameter:

- **Free Disk:** Add the unmounted or unpartitioned disks on image storage.



Note:

Configuring a free disk will format the selected disk, completely clearing all partitions, file systems, and data on the disk.

- **Mount Path:** Absolute path for mounting storage on image storage.



Note:

Do not use system directories such as `/`, `/dev/`, `/proc/`, `/sys/`, `/usr/bin`, `/bin`. Using system directories might cause image storage unable to work properly.

- **Retrieve Existing Images:** Default: disabled. If enabled, retrieves existing image files from the mount path in this image storage.
- **Image Sync Network:** The network CIDR used for image synchronization between standalone image storage within the same management node. If not specified, the management network is used by default. If both the source and target image storage set this parameter, only the image sync network of the target image storage takes effect.
- **Data Network:** The network CIDR used for data communication between compute nodes and the image storage. If not specified, the management network is used by default.

b) Configure advanced settings.

- **Reserved Capacity for Image Storage:** Default: 1 GB. The capacity reserved for use by the image storage.
- **Image Blob Upload Concurrency:** Default: 1. Specify the blob upload concurrency when you upload an image. Valid range: 1 to 16.

- **Image Blob Download Concurrency:** Default: 1. Specify the blob download concurrency when you download an image. Valid range: 1 to 16.

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

9.4.1.2 Add a Distributed Image Storage

A distributed image storage stores image files through 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, set the following parameters:
 - a) Set basic information.
 - **Name:** Image storage name.
 - **Description:** Image storage description.
 - **Type:** Displays distributed image storage.
 - **Data Center:** Current data center.
 - b) Set image storage configurations.
 - **Monitoring Node:** Add a monitoring node by specifying the monitoring node IP, SSH port, username, and password.
 - **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, a default image storage pool is created automatically.
 - **Data Network:** The network CIDR used for data communication between compute nodes and the image storage. If left blank, the management network is used by default.
 - c) Configure advanced settings.
 - **Reserved Capacity for Image Storage:** Default: 1 GB. The capacity reserved for use by the image storage.
 - **Image Blob Upload Concurrency:** Default: 1. Specify the blob upload concurrency when you upload an image. Valid range: 1 to 16.
 - **Image Blob Download Concurrency:** Default: 1. Specify the blob download concurrency when you download an image. Valid range: 1 to 16.

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

9.4.1.3 Manage an Image Storage

Modify an Image Storage

1. Navigate to **Menu > Image Storage**.
2. Right-click the image storage and select the action you want to perform from the list of actions that appears.
 - If you need to edit the name and description of the image storage, select **Edit Name and Description**.
 - If you need to modify the configuration information of the image storage, select **Modify Configuration**.
 - If you need to update the password for the image storage, select **Update Password**. After updating, you need to manually reconnect the image storage for the changes to take effect.

View Image Storage Usage

If you need to understand the storage usage information of the image storage, you can navigate to the image storage detail page to view the usage information. You can also clean up invalid data that has been permanently deleted and expired temporary data in the image storage based on your actual situation by clicking **Actions > Data Cleanup**.



Note:

Deleting an image file completely and deleting virtual machines created using that image completely allows you to execute a data cleanup operation to release storage space on the image storage. During data cleanup, please avoid performing data write-related operations.

Delete an Image Storage

If you need to delete an existing image storage, on the target image storage page, click **Actions > Delete** to delete it. You can also delete image storage in bulk on the data center **Image Storage** sub-page.

9.4.2 Image Basic Operations

9.4.2.1 Add an Image

You can add various formats of system images or disk images to the image storage. When creating a virtual machine for the first time, the system will download the image to the data storage

as an image cache. Using a URL or by uploading locally, you can add QCOW2, ISO, RAW, and VMDK image files to the image storage.

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, set the following parameters:
 - **Image Storage:** The current image storage
 - **Name:** Image name
 - **Description:** Image description
 - **Image Type:** Image type, including system images and disk images
 - **Image Format:** Different types of images support different image formats, including *.qcow2*, *.iso*, *.raw*, and *.vmdk* formats
 - **Image Path:** Add an image through a URL path or by uploading a local file
 - **URL:** Supports adding an image via HTTP/HTTPS/FTP/SFTP format or through an absolute path on a single-node image repository
 - **Upload File:** Choose a file that matches the selected image format and is accessible by the current browser to upload directly
4. Review the configuration and click **OK**.

9.4.2.2 Synchronize Images

Context

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

By default, image synchronization between image storage uses the management network. You can set a dedicated image sync network when adding an image storage or modifying image storage settings to reduce management network load. When both source and target image storage have set image sync networks, only the image sync network of the target image storage takes effect during the actual 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 target images.
5. Click **Bulk Action > Synchronize Image**.
6. In the **Synchronize Image** dialog, select the target standalone image storage.
7. Click **OK**.

9.4.2.3 Migrate an Image

Prerequisites

- You can migrate images between distributed image storage.
- Before migrating images, ensure the network connectivity between the monitoring nodes of the two distributed image storage.
- When migrating an ISO image across data centers, virtual machines using the ISO image need to detach the ISO image before they can start normally.

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 the image storage. You can manually clean up the data in the image storage. Once cleaned up, the data cannot be recovered. Proceed with caution.

9.4.2.4 Delete an Image

Prerequisites

(Optional) The platform provides deletion protection for images. You can define how resources are deleted by customizing the deletion policy in system parameters. By default, the platform adopts a delayed deletion for Image deletion policy (retained for 3 days). Deleted resources are first moved to the recycle bin and permanently deleted after the retention period. For more information, see [System Parameters](#).

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 multiple images, go to the image storage's image tab, select the VMs you want to delete, and then click **Bulk Action > Move to Recycle Bin**.

The delete button label changes based on the **Image Deletion Policy**. When the deletion policy is set to immediate deletion, the button appears as "Delete". When the deletion policy is set to delayed deletion or never delete, the button appears as "Move to Recycle Bin".

4. After acknowledging the risk, click **OK**.

9.4.3 Image Cross-Platform Usage

9.4.3.1 Export Image

Select an image and export it. The exported image will appear in the **Export List** tab of the root node. Once an image has been exported, it cannot be exported again.

Procedure

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



Note:

- For large image files, it is recommended to select export only. After exporting, you can go to the export list page to view and download.
- For medium and small image files, you can directly choose to export and download.

9.4.3.2 Export VM as OVA Template

Export a virtual machine in the .ova file format to a standalone image storage within the same data center. The exported OVA template does not include ISO drives or GPU devices.

Prerequisites

Before exporting an OVA template, ensure that the virtual machine is powered off.

Procedure

1. In the navigation pane, choose **Inventory > Host and VM**.
2. Right-click the target virtual machine, then select **Template > Export OVA Template**.


3. In the **Export OVA Template** dialog, choose the export method and image storage, and then click **OK**.

**Note:**

- For large VM files, it is recommended to select export only. After exporting, you can go to the export list page to view and download.
- For medium and small VM files, you can directly choose to export and download.

4. After the export completes, click **Inventory > Root Node**.

5. Click **Export List** to view the exported VM OVA template.

Figure 9-1: View the Exported VM OVA Template

9.5 Virtual Machine

A virtual machine is a core computing resource of the platform. Once you have prepared the necessary physical computing, storage, and networking resources and environment for the virtual machine, you can create a new virtual machine and manage its entire lifecycle according to your business scenarios.

9.5.1 Virtual Machine Basic Operations

9.5.1.1 Create a New Virtual Machine

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

Prerequisites

- The platform provides system parameters for VMs to control default settings globally. Before creating a VM, you can customize VM-related settings in the system parameters to control the default VM features. For more information, see [System Parameters](#).
- The platform provides default capacity and feature configurations for VMs, such as **CPU**, **Memory**, **Disk 1**, **NIC 1**, **Other Hardware**, and **Advanced Settings**. You can quickly create a virtual machine based on the default platform configurations.
- You can create VMs without an OS image. You can add the disk 1 by creating a new disk. A CD/DVD drive is optional.
- You can create VMs without NICs and configure NICs and their feature later.

Procedure

1. In the left 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, set the following parameters to complete the VM basic information:
 - **Name**: Virtual machine name.
 - **Quantity**: Number of virtual machines to create.
 - **Group**: VM group. Uses default group if not specified.
 - **Location**: Host or cluster where the VM resides.
 - **OS**: VM operating system. Supports mainstream OS including Linux and Windows.
 - **HA**: Automatic restart mechanism after abnormal shutdown. For more information, see [VM HA](#).
 - **Power Status**: Whether to power on the VM after creation. Default: enabled.
4. Set the following parameters to complete the VM hardware configurations:
 - a) CPU: Customize the VM CPU cores and features.
 - **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 VM CPU model with host CPU model to inherit host CPU features. Default: None. For more information, see [VM CPU](#)
 - **CPU Resource Priority**: For critical VMs, set to High to gain higher CPU contention capability when host is overloaded. For more information, see [VM Resource Contention](#).

- **CPU Clock Speed Limit:** Set the upper limit of host CPU resources that a VM can occupy. Valid range: 1% to 100%. 100% or blank means no limit.
- **CPU NUMA Binding:** Bind the virtual machine's CPU (vCPU) to the host's CPU (pCPU) to improve its performance. For more information, see [VM Resource Contention](#).
- **CPU Hot Plug:** Whether to support online CPU modification. Default: enabled. CPU hot-plugging and memory hot-plugging must be enabled or disabled together.

**Note:**

Only some operating systems support hot-plugging. You can click **View** to check supported OS list.

- **CPU Hypervisor Tag:** Whether to enable the virtualization (hypervisor) flag. If disabled, skips the VM virtualization environment detection by applications. Default: enabled.

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

- **Memory:** VM memory capacity.
- **Memory Resource Priority:** For critical VMs, set to High to gain higher memory contention capability when host is overloaded. For more information, see [VM Resource Contention](#).
- **Memory Hot Plug:** Whether to support online memory modification. Default: enabled. CPU hot-plugging and memory hot-plugging must be enabled or disabled together.

**Note:**

Only some operating systems support hot-plugging. You can click **View** to check supported OS list.

c) Disk 1: Customize the storage location, capacity, and features of the virtual machine's system disk. You can create a new disk or choose an existing system image.

- **Storage Location:** Data storage location for VM disk. Uses automatic allocation if not specified.
- **Capacity:** System disk capacity when creating a new disk.
- **System Image:** Select a system image when using system image. Supported format: raw and qcow2.
- **Bus Type:** Specify the bus type for a virtual machine's disk. Options include Virtio, IDE, Virtio SCSI, and SCSI. By default, Disk 1 uses Virtio bus type on Linux systems, and IDE bus type on Windows and Other systems.

- **Provision Method:** Choose how to allocate disk storage space when using SAN storage. Default: thin provisioning.
 - Thin Provisioning: Allocates storage space based on actual usage, achieving higher storage utilization.
 - Thick Provisioning: Pre-allocates the required storage space when creating the disk, providing sufficient storage capacity and ensuring storage performance.
- **Caching Mode:** Whether to use host page cache for write operation and if used, whether the data is written to the storage device before returning success. Default: none.
- **AIO Acceleration:** Whether to enable asynchronous I/O (AIO) acceleration in the VM kernel. Default: disabled.

**Note:**

To enable AIO acceleration, make sure the cache mode is set to none.

- **QoS:** Whether to set read/write bandwidth and IOPS limits. Default: no limit. For more information, see [VM QoS](#).

Disks other than Disk 1 are data disks, such as Disk 2. Add data disks by clicking **Add Hardware > Disk**. A single virtual machine can support up to 24 disks (including Disk 1).

- Data disk addition methods: new disk, disk image, existing disk, and RDM disk.
- By default, data disks use Virtio bus type on Linux and Windows systems, while only IDE is supported for Other systems.
- When the storage location is ZHPS distributed storage, the data disk's bus type only supports Virtio.
- You can share a disk when the storage location is nSDS distributed storage and the bus type is Virtio SCSI.
- You can share a disk when the storage location is SAN storage, the bus type is Virtio SCSI, and the provision method is thick provision.
- You cannot modify the bus type or QoS of a shared disk.

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

- **NIC Model:** Set the NIC model. Supported models: e1000, rtl8139, virtio, and SR-IOV.
- **Port Group:** The port group of the distributed switch for the VM NIC.
- **State:** Whether to automatically enable NIC when VM powers on.

- **NIC Queue Number:** Use multiple queues to send and receive network packets to improve network PPS and bandwidth performance.
- **MAC Address:** Specify a MAC address. Default: Auto Generated.
- **Specify IP Address:** Specify an IP address. Default: automatically assigned.

If the selected distributed port group has DHCP service disabled, you can use VMTools to specify an IP address for the virtual machine. For more information, see [Virtual Machine VMTools](#).

- **Assign DNS:** Specify a DNS address. Default: Auto Allocated.
- **Security Group:** Associate security groups with the VM NIC to control east-west traffic. The smaller the number on the left side of the associated security group, the higher the priority for taking effect. For more information, see [Security Group](#)



Note:

Configure carefully to avoid rule conflicts between security groups.

- **QoS:** Whether to set bandwidth limits on packet transmission for the VM NIC. Default: no limit. For more information, see [VM QoS](#).

Add multiple NICs by clicking **Add Hardware > NIC**.

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

Add multiple CD/DVD drives by clicking **Add Hardware > CD/DVD Drive**. You can add a maximum of three CD/DVD drives to a virtual machine.

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

Supports physical GPUs and vGPUs. Add multiple GPU devices by clicking **Add Hardware > GPU Device**.

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

Supports direct connection and redirection. Add multiple USB devices by clicking **Add Hardware > USB Device**. You can add a maximum of one USB device to a virtual machine.

- h) PCIe Device 1: Attach a PCIe device for the virtual machine.

You can load a PCIe device that has been switched to pass-through mode on the host. To add a PCIe device, click **Add Hardware > PCIe Device**.

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

You must first add an available key provider and set the virtual machine boot mode to UEFI. To add a TPM device, click **Add Hardware > TPM**. A single virtual machine supports a maximum of one TPM device.

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

- **Graphics Card Type:** Specify the default graphics card type when powering on a virtual machine. Supports vga, virtio, qxl, and cirrus, providing basic/high-definition/high-performance video functionality experiences. Default: vga for x86 VMs and virtio only for ARM VMs.
- **Total Graphics Memory:** Fixed at 16 MB for vga or cirrus. Configurable for qxl.
- **Audio Card Type:** Specify the default audio card type when powering on a virtual machine. Supports HDA (ICH6), HDA (ICH9), and AC97. Default: HDA (ICH6).
- **Motherboard Type:** Specify the default motherboard type when powering on a virtual machine. Supports i440fx and q35. Default: i440fx.

5. Set the following parameters to complete the VM advanced settings:

a) Configure general options.

- **Tag:** Attach tags to identify VMs. For more information, see [Tag Management](#).
- **Hostname:** VM hostname.
- **VM Scheduling Group:** Join a VM scheduling group for host allocation based on scheduling policies associated with the scheduling group. For more information, see [VM Scheduling Policy](#).
- **Sync with Host BIOS:** Whether to synchronize the Windows VM BIOS clock with the host's BIOS clock. Default: disabled. For more information, see [VM Time Synchronization](#).
- **User Data:** User-defined data. Upload custom parameters or scripts to perform custom configurations or specific tasks on the virtual machine. For more information, see [VM User Data](#).

b) Configure remote access settings.

- **Console Mode:** VM console mode. Options include VNC (default), SPICE, VNC+SPICE.
- **Console Password:** VM console password. Supports manual input or random generation. Allowed characters: letters, numbers, and the following special characters: - `=[];',./~!@#\$%^&*()_+|{}:"<>?

You can specify whether to enforce setting a console password and password strength through VNC Console Password in the security settings. If enabled, you need to set the console password according to the specified strength requirements when creating a virtual machine.

- **USB Redirection:** Redirects USB devices from the VDI client to the VM. Default: disabled.

c) Configure login authentication settings.

- **None:** Do not set login password or SSH key.
- **Password:** System login password for Linux VM (root) or Windows VM (administrator). Supports manual input or random generation. Allowed characters: letters, numbers, and the following special characters: -`=[];',./~!@#\$\$%^&*()_+|{}:"'<>?



Note:

Before setting the password, ensure that the virtual machine image has `cloud-init` installed.

- **SSH Key:** Inject SSH Key for password-free login to Linux VM.



Note:

Before injecting the SSH Key, ensure that the virtual machine image has `cloud-init` installed.

d) Configure VMTools.

- **Failure Response Policy:** Set an automatic response action for VM failures (Windows BSOD or Linux guest hang). Options include No Action, Reboot, and Shut Down.
- **Time Synchronization:** Whether to automatically synchronize VM time with host system time. Default: enabled.

e) Configure boot options.

- **Boot Order:** OS boot priority sequence. Options include hard disk, CD-ROM, and network. By default, the BIOS boots from the hard disk, and if no boot device is found, it cannot load the system.
- **BIOS Mode:** The BIOS boot mode supports Legacy and UEFI. Legacy is the default in x86 clusters, while UEFI is the default in ARM clusters.
- **Secure Boot:** You can enable this option only when the BIOS mode is set to UEFI. When enabled, the system validates the digital signature of all boot components (such as

boot loaders, OS Kernels) during VM startup. Only trusted software with valid signatures can be loaded.

- **BIOS Post Delay:** BIOS screen timeout duration. Default: 10 seconds.

f) Configure other options.

- **Hide KVM Virtualization Flag:** Controls CPU virtualization flag. When enabled, inserts `<hidden state="on">` into the `<kvm>` element in the newly started VM XML. Default: disabled.
- **VMware I/O Port Simulation:** Whether to allow a KVM virtual machine to emulate the I/O ports in a VMware virtualization environment, making the KVM virtual machine compatible with VMware I/O port standards. Default: disabled.

The main purposes of this option:

- **Migration and Compatibility:** Allows migrating virtual machines from a VMware environment to a KVM environment, or running both VMware and KVM virtual machines in a mixed environment without significant configuration changes.
- **Testing and Development:** Developers and testers can use KVM virtual machines to simulate a VMware environment to test applications or configurations in a VMware-like environment without requiring actual VMware licenses.
- **Anti-Spoofing Mode:** Whether to enable anti-IP/MAC spoofing and ARP deception features. When enabled, the virtual machine can only communicate with the outside world using the IP/MAC address allocated by the platform. Default: disabled.
- **Cross-Cluster HA Policy:** Whether VM HA migration supports cross-cluster.



Note:

This policy only affects automatic VM migrations in shared storage scenarios (distributed storage, NFS storage, SAN storage), such as VM migration in host maintenance or host replacement. This policy does not affect other actions such as manual hot migration of virtual machines (changing the host), specifying a host to start a VM, or dynamic resource scheduling (DRS) strategies that change the host.

- **Hyper-V:** Whether to enable Hyper-V emulation for the VM for nested virtualization scenarios in Windows systems. Default: disabled.
- **EmulatorPin:** Whether to assign all other threads than vCPU threads and IO threads of a VM to host physical CPUs (pCPUs). You can assign by NUMA nodes. Default:

unassigned, which is virtual machine-related threads run on corresponding pCPUs according to system scheduling.

- **Auto-Converge:** Whether to enable auto-convergence mode for KVM virtual machine hot migrations. Default: disabled.

If the virtual machine remains under high business load for a long time and the application is moderately sensitive to performance, it is recommended to enable auto-convergence mode to improve migration success rates.

- **PCI Hot Plug:** Whether to allow hot plugging of PCI devices in the virtual machine. Default: enabled.



Note:

If hardware compatibility errors occur during hot plugging or if the hardware device is not supported, you can disable this switch.

- **CPU Vendor ID:** If the virtual machine is running on a host with Hygon CPUs, it is recommended to set the VM's CPU vendor ID to AuthenticAMD to ensure compatibility with various operating systems and maintain normal VM operation. If set to None, certain operating systems may experience compatibility issues.



Note:

This parameter is hidden when the location of the virtual machine is set to either cluster or auto-allocated.

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

What's next

Some VM configurations require VMTools. After VM creation, it is recommended to install VMTools to enable certain configurations. For more information about VMTools, see [Virtual Machine VMTools](#).

9.5.1.2 Import a Virtual Machine

Upload an OVF file to quickly import virtual machines, facilitating VM migration across different platforms.

Prerequisites

- The platform provides system parameters for VMs to control default settings globally. Before creating a VM, you can customize VM-related settings in the system parameters to control the default VM features. For more information, see [System Parameters](#).

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 Virtual Machine** dialog, set the following parameters:

Upload Information

- **Image Storage:** Temporary storage for the uploaded template file. Supports single-node image repositories and distributed image repositories. The template file will be automatically deleted after the virtual machine is created.
- **Template Type:** Supports uploading OVF type template files.
- **OVF File:** Upload an OVF format file. Only a single file can be uploaded.
- **VMDK File:** Upload VMDK format files from the OVF template. Must be consistent with the file configuration defined in the OVF format file, including the number of files and configuration details.
- **MF File:** Upload an MF format file from the OVF template. Only a single file can be uploaded.

Basic Information

- **Name:** The name of the virtual machine.
 - **Group:** The group where the virtual machine resides. If not set, the default group will be used.
 - **Location:** The host or cluster location where the virtual machine is running.
 - **Storage Location:** The data storage location for the virtual machine's hard disk. Supports local storage, NFS storage, SAN storage, as well as distributed storage. If not set, automatic allocation will be used.
 - **HA:** The auto-restart mechanism for the virtual machine in case of an abnormal shutdown. It is disabled by default. For more information, see [VM HA](#).
 - **Power Status:** Whether the virtual machine should automatically power on after creation. By default, it starts automatically.
4. Review the configuration and click **OK**.

What's next

Some VM configurations require VMTools. After VM creation, it is recommended to install VMTools to enable certain configurations. For more information about VMTools, see [Virtual Machine VMTools](#).

9.5.1.3 Create a Virtual Machine from a Template

Create identical virtual machines from a template. You'll be able to customize hardware, software, and other configurations.

Prerequisites

- The platform provides system parameters for VMs to control default settings globally. Before creating a VM, you can customize VM-related settings in the system parameters to control the default VM features. For more information, see [System Parameters](#).
- Before creating a virtual machine based on a template, make sure there is already a virtual machine template available in the platform.

Procedure

1. In the navigation pane, choose **Inventory > VM and Host**.
2. Right-click a valid parent resource (cluster or host) of virtual machines and select **New Virtual Machine**.
3. In the **Select VM Creation Type** dialog, choose **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, set the following parameters:
 - a) Complete the template information configuration.
 - **Template:** Select a template and create a virtual machine based on this template.
 - b) Complete the basic information configurations.
 - **Name:** The name of the virtual machine
 - **Quantity:** The number of virtual machines to create this time



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:** The group the virtual machine belongs to; if not set, the default group will be used
- **Location:** The host or cluster location of the virtual machine
- **OS:** The operating system of the virtual machine, supporting mainstream operating systems including Linux and Windows
- **HA:** Automatic restart mechanism for the virtual machine after abnormal shutdown, default is disabled. For more information, see [VM HA](#)
- **Power Status:** Whether the virtual machine automatically powers on after creation, default is to power on automatically

c) Complete the hardware information configurations.

- **CPU:** Supports adjusting the total number of cores and the number of cores per socket, and setting CPU hot plug.
- **Memory:** Supports adjusting the memory size and setting memory resource priority.
- **Disk:** Supports modifying the cache mode of the disk and setting AIO acceleration. You can add a new hard disk to the virtual machine by clicking **Add Hardware > Disk**. The new hard disk allows customization of its capacity and properties.
- **NIC:** Supports modifying the NIC model, port group, NIC queue number, MAC address, IP address, and DNS assignment. You can add a new NIC to the virtual machine by clicking **Add Hardware > NIC**. The new NIC allows customization of its address and properties.
- **CD/DVD Drive:** Supports loading ISO image files onto the virtual machine for booting from an ISO optical drive.
- **GPU Device 1:** Loads a GPU device onto the virtual machine, supporting physical GPU devices and vGPU devices.

You can add a GPU device by clicking **Add Hardware > GPU Device**.

- **USB Device 1:** Loads a USB device onto the virtual machine, supporting direct connection and redirection.

You can add a USB device by clicking **Add Hardware > USB Device**. A single virtual machine supports adding up to 1 USB device.

- **Other Hardware:** Does not support modifying the graphics card and sound card configuration in the template file.

d) Complete the advanced settings.

General Options

- **Description:** Displays the description recorded in the template. You can modify the VM description.
- **Tag:** Displays the tag recorded in the template. You can customize different 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 VM configurations require VMTools. After VM creation, it is recommended to install VMTools to enable certain configurations. For more information about VMTools, see [Virtual Machine VMTools](#).

9.5.1.4 Register a Virtual Machine

You can register a virtual machine configuration file and data file to create a corresponding virtual machine record on the platform. After registration, the virtual machine is managed uniformly by the platform.

Prerequisites

- Make sure the platform is installed with a valid Advanced Edition license.
- Only SAN storage, local storage, and NFS storage support virtual machine registration. Before registration, the data storage must be attached to the cluster, and the cluster must contain a host.
- The same VM configuration file cannot be registered multiple times on the same host.
- For VMs with TPM, ensure 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, set the following parameters:

- **Configuration File:** Select the configuration file for the virtual machine you want to register.
- **VM Name:** Set a name for the virtual machine.
- **Group:** Select the group for the virtual machine.
- **Location:** Specify where the virtual machine runs.

4. Click **OK**.

9.5.1.5 Clone a Virtual Machine

If you want to reuse the configuration of a virtual machine, you can clone the virtual machine.

NexaSphere provides full VM cloning with two methods: full cloning and instant full cloning

Prerequisites

- The source virtual machine must be in a running, paused, or shut down state.



Note:

Cloning running VMs only clones the data that has already been written to disk at the start of cloning and does not include real-time cache data.

- To ensure data integrity, it is recommended to pause or power off high I/O virtual machines before the cloning.
- Detach all shared disks from the source virtual machine. Cloning of virtual machines with shared disks is currently unsupported.
- To avoid login errors, do not set a static IP address on the source virtual machine.
- The cluster where the virtual machine resides has sufficient compute, storage, and network resources.

Procedure

1. On the target VM page, click **Actions > Clone > Clone Virtual Machine**.

2. In the **Clone Virtual Machine** dialog, set the following parameters:

- **Name:** Set the name for the cloned virtual machine.
- **Quantity:** Set the number of virtual machines to be cloned.

When cloning multiple VMs, the system automatically appends suffixes "-1/-2/-3" to distinguish between cloned resources.

- **Clone Method:** Choose a type of cloning method.

- **Full Clone:** The cloned VM is independent of the source VM, and the performance is completely unaffected after cloning, but the VM starts slow.
- **Instant Full Clone:** The cloned VM starts quickly, the VM is eventually independent of the source VM, and the performance is completely unaffected after cloning.

**Note:**

When you use Instant Full Clone to clone a VM, the system automatically performs Flatten to eventually achieve data independence. During flattening, operations on VMs/disks will be conducted after the flattening is completed.

- **Data Storage:** By default, the system automatically assigns data storage for cloned virtual machines, but you can manually specify a data storage.
 - Auto Allocated: The cloned VM and its disks will use the same data storage as the source VM and its disks.
 - Manual Allocation: The cloned VM and its disks will use the data storage you specify.

**Note:**

1. You can only specify data storage when using full clone.
2. If Disk 1 of the source VM uses local storage while the other disks use different types of data storage, the clone operation will fail if local storage is unavailable.
3. When specifying SAN storage as the data storage, you can set separate provisioning methods for Disk 1 and other disks.

- **VM Scheduling Group:** Add cloned VMs to a virtual machine scheduling group. The virtual machine will be scheduled based on the scheduling policy associated with group. A virtual machine can join only one VM scheduling group. For more information, see [VM Scheduling Policy](#).
- **Power Status:** Select whether to automatically power on virtual machines after the cloning.

What's next

The cloning operation duplicates the source virtual machine's configuration, installed applications, and credentials to the new cloned VM. However, the cloning operation does not automatically replicate all associated settings. You must manually configure additional parameters for the cloned VM as needed. For example, to maintain consistency with the source VM, you may need to manually assign identical tags to the cloned VM.

If the source virtual machine has a console password configured, you can restart the cloned virtual machine to activate the password setting.

9.5.1.6 Access a Virtual Machine

You can access a virtual machine by launching the VM console with one click. In the VM console, you can perform various actions, such as installing the operating system, configuring the system, executing commands, and running applications. NexaSphere provides these access options:

- Access a VM by using the console
- Access a VM by using SSH

9.5.1.6.1 (Optional) Manage VM Access and Boot Options

Before accessing a virtual machine, you may configure the following settings according to your business needs: boot options, remote access, and login authentication.

Manage Boot Options

The boot options include the following settings:

- **Boot Order:** Define the sequence for loading the operating system during VM startup.

For example, if you set the boot order as (1) Disk 1, (2) CD/DVD Drive, and (3) Network, the VM boots in this sequence:

1. First the VM attempts to boot from disk 1. If successful, the VM will not try CD/DVD drive.
2. If booting from disk 1 fails, the VM attempts to boot from CD/DVD drive. If successful, the VM will not try the network.
3. If booting from 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 will not start.

- **BIOS Mode:** Select the BIOS mode based on the image format:
 - **Legacy:** Supports x86 architecture and all operating systems.
 - **UEFI:** Required for aarch64 architecture. Supports Windows and CentOS. Windows 7/2008 requires compatibility support module (CSM).

Mismatched BIOS mode may cause virtual machines to malfunction:

- For qcow2 or raw images, select the BIOS mode that is consistent with the template.
- For iso images, you can select a BIOS mode as needed. Then, the system will be booted accordingly.

- **BIOS Post Delay:** Set the automatic delay time for BIOS interface. Valid value: 1 to 60 seconds . If you do not perform any operation during this period, the system automatically proceeds to boot.

Prerequisites

To configure boot options, make sure the virtual machine is shut down.

Procedure

1. On the target virtual machine page, click **Advanced Settings**.
2. On the **Advanced Settings** tab, choose **Boot Options**.
3. Click **Edit**, then configure boot order, BIOS mode, BIOS post delay as needed.

Manage Remote Access Settings

The remote access includes the following settings:

- **Console Mode:** Select the protocol type for connecting to the VM console:
 - **VNC:** Support both Linux and Windows systems. Primarily used for Linux server management with typical network traffic around 100 KB.
 - **SPICE:** Support Linux systems with superior color/audio/video/USB capabilities. Ideal for virtual desktop applications with typical network traffic around 10 to 20 MB. In this mode, you can configure the number of connected displays and video streaming mode (off, all, or filter).
 - **VNC+SPICE:** Support both VNC and SPICE protocols. You can configure the number of connected displays and video streaming mode (off, all, or filter).
- **Console Password:** Set the console access password.
 - **Character requirements:** Support letters, numbers, and these special characters: -`=[];','./~!@#\$\$%^&*()_+|{}:."<>?
 - **Length requirements:** 6 to 8 characters.

You can customize the VNC console password strength. For more information, see [Security Settings](#).

Prerequisites

To configure remote access, make sure the virtual machine is shut down.

Procedure

1. On the target virtual machine page, click **Advanced Settings**.

2. On the **Advanced Settings** tab, choose **Remote Access**.
3. Click **Edit**, then configure console mode and console password as needed.

Manage Login Authentication Settings

The login authentication includes the following settings:

- Set SSH Key:
 1. After installing cloud-init, the SSH authentication is disabled by default. To enable it, set the `ssh_pwauth` parameter to **1** in `/etc/cloud/cloud.cfg`.
 2. Generate an SSH key pair using the `ssh-keygen` command. 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 injecting SSH KEY during new VM creation, it takes effect after the first boot.
 - For existing VMs receiving SSH KEY for the first time, you must reboot the VM.
 - To re-inject SSH KEY for VMs with existing configurations, first run `rm -rf /var/lib/cloud/instances` to clear previous settings before injecting the new SSH KEY and rebooting.
 - Deleting SSH KEY only removes the record from the system, not from VM configurations. To completely remove SSH KEY, manually delete it from `/root/.ssh/authorized_keys` in the VM.
- Change VM Password:
 1. Modify the login name as needed. Options include system default and custom login name.
Default login name: root for Linux and Administrator for Windows.
 2. Set a new password.
 - Character requirements: Support letters, numbers, and these special characters: `-`=[]',./~!@#%&*()_+|{}:"<>?`

You can customize the VM password strength. For more information, see [Security Settings](#).

Prerequisites

- Before configuring SSH public key, you must install VMTools (cloud-init). Recommended versions: 0.7.9, 17.1, 19.4, or later.

- Before changing VM password, ensure the VM is running and has VMTools (QGA) installed.

Procedure

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

9.5.1.6.2 Access a VM by using Console

You can quickly access a virtual machine through the console to perform various actions, such as install the operating system, configuring system, executing commands, and running applications.

Prerequisites

- The virtual machine must be in the running state.
- (Optional) Configure console mode and console password as needed. For more information, see [Manage Remote Access Settings](#).
- (Optional) Set up a proxy address to access the VM console. For more information, see [Console Proxy Management](#).

Procedure

1. In the navigation pane, choose **Inventory > VM and Host**.
2. Choose a target virtual machine 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. The console provides quick-access buttons in the left sidebar inside the console window.
 - **Local Command Paster**: Click this button and the paste command dialog appears. You can paste commands here and click **OK** to run the commands.
 - **Tools**: Provide a collection of hotkeys, including Ctrl, Alt, Win, Tab, Esc, and Ctrl+Alt+Del.
 - **Power Management**: Allow you to manage the VM power status. Actions include shut down, reboot, pause, resume, and power off.
 - **Setting**: Choose whether to enable a read-only mode in the console. When enabled, you cannot enter commands or perform other actions in the console.

9.5.1.6.3 Access a VM by using SSH

You can access a virtual machine by using the SSH connection protocol.

Prerequisites

- You have network connection and root access privileges to the target VM.
- (Optional) Set SSH Key for password-free login and change VM password as needed. For more information, see [Manage Login Authentication Settings](#).

Procedure

1. Log in to the virtual machine directly using a remote login software on your local computer.

```
# 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 using an SSH private key, first add the SSH public key to the target virtual machine, then run the following command:

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

`${private_key.pem}` is the path to your private key file, `${UserName}` is your login username, and `${IpAddress}` is the target VM IP address.

9.5.1.7 Modify a Virtual Machine

If you have already created one or more virtual machines, you can modify the virtual machine configurations according to your business scenarios as needed.

Before Modifying Virtual Machine Configurations

Before modifying the following configurations, you need to install VMTools on the virtual machine:

- Modify failure response policy
- Change the virtual machine password
- Modify the hostname for Windows virtual machines
- Modify time synchronization

To modify the following configurations, you need to shut down the virtual machine:

Module	Modification Item
Operating System	System Image
CPU and Memory	Number of Cores per Socket

Module	Modification Item
	CPU Hot Plugging & Memory Hot Plugging
	CPU Mode
	CPU Affinity
Network Adapter	Network Adapter Multi-Queue Count
Graphics Card	Graphics Card Type
	Total Graphics Memory
General Options	Hostname
Login Options	Password
Remote Access	Console Mode
	Console Password
	USB Redirection
	SSH Key
Boot Options	BIOS Mode
	Boot Order

Modify Virtual Machine Configuration

Single Operation:

- If you only need to modify the virtual machine name and description, you can click on the **Actions > Edit Name and Description** in the target virtual machine page to make the changes.
- If you need to modify the [Basic Information Settings](#) and [Hardware Information Settings](#), you can click on the **Modify Configuration** operation or the **VM Hardware** settings on the **Overview** page to make the changes.
- If you need to modify the [Advanced Settings](#), you can click on the **Actions > Advanced Settings** in the target virtual machine page and modify the corresponding settings as needed.
- If you need to modify the virtual machine system settings, such as resetting the virtual machine, changing the group the virtual machine is located in, issuing network configurations, or changing the owner of the virtual machine, you can click on the **Actions > System Settings** in the target virtual machine page and modify the corresponding settings as needed.
- If you need to modify the tags bound to the virtual machine, you can click on the **Actions > Tag Management** in the target virtual machine page to make the changes.

Bulk Operations:

- If you need to modify the system settings of multiple virtual machines, such as resetting the virtual machines, changing the groups they are located in, or changing the owners of the virtual machines, you can do so on the child virtual machine page corresponding to the parent/grandparent resource (host, cluster, data center) of the target virtual machine. Select the target virtual machines and click on the **Bulk Actions > System Settings** and modify the corresponding settings as needed.
- If you need to modify the tags bound to multiple virtual machines, you can do so on the child virtual machine page corresponding to the parent/grandparent resource (host, cluster, data center) of the target virtual machine. Select the target virtual machines and click on the **Bulk Actions > Tag Management** to make the changes.

9.5.1.8 Delete a Virtual Machine

Prerequisites

- You cannot move virtual machines to recycle bin or delete virtual machines that are in the running, paused, crashed, or unknown state. Shut down the VM first before perform the action.
- (Optional) The platform provides deletion protection for virtual machines and disks. You can define how resources are deleted by customizing the deletion policy in system parameters. By default, the platform adopts a delayed deletion for VM deletion policy (retained for 7 days) and disk deletion policy (retained for 3 days). Deleted resources are first moved to the recycle bin and permanently deleted after the retention period. For more information, see [System Parameters](#).

Procedure

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

To delete multiple virtual machines, go to the virtual machine tab of the parent resources (host, cluster, data center). For example, on the host's **Virtual Machine** tab, select the VMs you want to delete, then click **Bulk Action > Move to Recycle Bin**.

The delete button label changes based on the **VM Deletion Policy**. When the deletion policy is set to immediate deletion, the button appears as "Delete". When the deletion policy is set to delayed deletion or never delete, the button appears as "Move to Recycle Bin".

4. In the confirmation dialog, select whether to simultaneously delete disks attached to the virtual machine (excluding shared disks).

When selected, disks will be deleted or moved to recycle bin according to the defined disk deletion policy.

5. After acknowledging the risk, click **OK**.

Result

The deletion releases the associated CPU, memory, and IP address resources.

9.5.2 VM Group Management

A VM group is a logical collection that users create based on business requirements to organize virtual machines. NexaSphere offers a default group; when you create new VMs, if you have not created your own groups or assigned them to a group, they will automatically be placed in the default group. This section explains how to create and use VM groups.

- [Create New VM Groups and Subgroups](#)
- [Add VMs to VM Groups or Subgroups](#)
- [Modify VM Group Name](#)
- [Delete VM Group](#)

Create New VM Groups and Subgroups

The platform provides multiple entry points for creating new VM groups. You can create a new VM group or subgroup from either of the following main entry points:

- In the left navigation pane, right-click the target data center and click **New VM Group**.
- In the left navigation pane, select the target data center. Then on the right side of the page, click **Actions > New VM Group**, or on the **Virtual Machine** tab under the **VM Group** tab, click **New VM Group**.

To create a new VM group, set the following parameters:

- **Name:** Name of the VM group
- **Description:** Description of the VM group

After you click **OK**, the creation is complete.

Once the VM group is created, if you need to create a subgroup within this group, you can do so in the resource section of the group in the left navigation pane or on the group's page on the right side of the platform. Click **Create Subgroup** and set the following parameters:

- **Affiliated to:** Shows the parent VM group of this subgroup
- **Name:** Name of the VM subgroup

Add VMs to VM Groups or Subgroups

After creating a VM group or subgroup, you can add VMs to it. The platform provides multiple entry points for adding VMs to a group. You can add VMs from any of the following main entry points:

- In the left navigation pane, right-click the target VM group and click **Create VM**. The newly created VM will belong to this group by default.
- In the left navigation pane, select the target VM group. Then on the right side of the page, click **Create VM**. The newly created VM will belong to this group by default.
- On the target VM's page, click **Actions > System Configuration > Change Group** to change the current VM's group to the target group.

Modify VM Group Name

If you need to change the name of a VM group, on the target VM group page, click **Actions > Edit Name** to modify the name.

Delete VM Group

If you no longer need a specific VM group, on the target VM group page, click **Actions > Delete** to remove the group. You can also delete multiple groups at once from the VMs tab under data center resources. After deleting a group, the VMs within it will be moved to the default group.

9.5.3 Virtual Machine VMTools

The Virtual Machine VMTools is a collection of drivers and tools that enrich virtual machine functionality and enhance performance.

VMTools Components

VMTools primarily includes the following tools and drivers, which may vary depending on the operating system.

OS	VMTools Components	Description
Linux	Advanced Monitoring Agent	An agent installed on virtual machines for collecting VM advanced monitoring data, including CPU, memory, and disk capacity data.

OS	VMTools Components	Description
	QEMU Guest Agent (QGA)	An application for interaction between virtual machines and hosts, without relying on network. It can read and deliver VM configurations, and push monitoring data. Note: Linux OS usually comes with built-in virtio drivers and cloud-init, so the VMTools for Linux does not include these two tools. If your Linux OS lacks them, you can install them manually through other methods.
Windows	Advanced Monitoring Agent	An agent installed on virtual machines for collecting VM advanced monitoring data, including CPU, memory, and disk capacity data.
	QGA	An application for interaction between virtual machines and hosts, without relying on network. It can read and deliver VM configurations, and push monitoring data.
	Virtio Driver	A collection of drivers to improve the VM performance. - SCSI controller driver: Used to improve the VM disk performance in a virtualization environment. - PCI simple communications controller driver: Used to realize the communications between a Windows virtual machine and the underlying KVM virtualization. - PCI device driver: Used to realize the balloon memory scaling. - Ethernet adapter driver: Used to improve the network performance of a Windows virtual machine in a virtualization environment.
	Cloudbase-init	A component helps realize the User Data import and other customized functions.
FreeBSD	Advanced Monitoring Agent	An agent installed on virtual machines for collecting VM advanced monitoring data, including CPU, memory, and disk capacity data.

Using VMTools

After installing VMTools on a virtual machine, you can achieve advanced monitoring, configuration issuance, and configuration reading.

- **Advanced Monitoring:** Advanced monitoring is the monitoring data obtained from the virtual machine by the advanced monitoring agent, which is periodically pushed to the host via DHCP service or QGA. You can use line charts to monitor real-time trends in virtual machine performance, including CPU, memory, and disk capacity.
- **Configuration Issuance:** Through QGA in VMTools, parameters set on the platform are actually issued to take effect inside the virtual machine, including the virtual machine hostname and network configurations.
 - **Hostname:** After installing the latest version of VMTools, users can modify the hostname of a running virtual machine, and the new hostname will be issued directly via QGA without needing to restart the virtual machine.

**Note:**

If you need to issue the hostname via QGA, please install VMTools before modifying the hostname.

- **Network Configuration:** When distributed port groups do not have DHCP services enabled, IP addresses, DNS, MTU, and other parameters set on the platform need to be issued via QGA to take effect on the NIC. After installing VMTools, the specified IP addresses and other parameters will automatically take effect.
- **Configuration Reading:** Through QGA in VMTools, the IP address configured on the virtual machine's NIC is automatically read, and the successfully read IP address can be displayed and managed on the platform.

**Note:**

- If the distributed port group to which the virtual machine's NIC belongs does not have DHCP enabled, and the virtual machine has VMTools installed and is in a running state, it supports automatically reading the virtual machine's NIC configuration information.
- If an IP address is already configured inside the virtual machine, it will be read and overwrite the IP address configured on the platform.
- If there is a duplicate IP address between the virtual machine's NIC IP address and the IP addresses in its port group, an alert will be triggered.

9.5.3.1 Install VMTools on Linux VMs

During the installation, you will be guided to first attach the VMTools ISO, then enter the VM console to manually execute the installation commands.

Prerequisites

- Make sure the virtual machine is in the running state.
- Make sure the virtual machine has installed Linux command tools, 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 attaching the VMTools ISO, the system automatically chooses a CD/DVD drive based on the following rules:

- When a CD/DVD drive is available, the system automatically selects the first one and attaches the VMTools ISO.
- When no CD/DVD drive is available, the system automatically detaches the original image on CD/DVD drive 1 and attach the VMTools ISO.

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

Click the button copies the following commands and enters the VM console:

```
# Create a mount point.
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. Enter the VM console. In the left toolbar, click **Local Command Paster**.
5. In the **Paste Command** dialog, paste the copied commands from the last step and click **OK** to run the commands.

What's next

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

9.5.3.2 Install VMTools on Windows VMs

During the installation, you will be guided to first attach the VMTools ISO, then enter the VM console to install VMTools following the prompted steps.

Prerequisites

Make sure the virtual machine is in the running state.

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 attaching the VMTools ISO, the system automatically chooses a CD/DVD drive based on the following rules:

- When a CD/DVD drive is available, the system automatically selects the first one and attaches the VMTools ISO.
- When no CD/DVD drive is available, the system automatically detaches the original image on CD/DVD drive 1 and attach the VMTools ISO.

3. Click **Launch Console**.

4. Enter the VM console and install VMTools following the prompted steps.

1. Load the VMTools image: Click install VMTools in the VMTools installation prompt to load the image to the virtual CD Drive.
2. Install VMTools: Run the VMTools installation program and install VMTools, commonly-used tools, and virtio driver.
3. Confirm the installation and reboot the virtual machine: Reboot the virtual machine to make the modification take effect after the installation.

What's next

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

9.5.3.3 VMTools Compatibility with OS

Different operating systems have varying levels of compatibility with VMTools. The following tables show the compatibility status between each operating system and VMTools.



Note:

In the table, **Yes** indicates that the component can be installed on the VM through VMTools, and its related functions can be used normally. **No** only means that the component cannot be installed on the VM through VMTools, or that dependent features may not work after installation. It does not imply that the component is unsupported—you can still install it through other methods.

Linux OS

OS	OS Release	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

OS	OS Release	Advanced Monitoring Agent	QGA
	Redhat Enterprise Linux Server 7.6 64-bit	Yes	Yes
Fedora	Fedora 30 64-bit	Yes	Yes
	Fedora 31 64-bit	Yes	Yes
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

OS	OS Release	Advanced Monitoring Agent	QGA
	SUSE Linux Enterprise Desktop 12 64-bit	Yes	Yes
	SUSE Linux Enterprise Desktop 15 64-bit	Yes	Yes
UOS	UOS V20 1050e	Yes	Yes
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 OS

OS	OS Release	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

FreeBSD OS

OS	OS Release	Advanced Monitoring Agent
FreeBSD	FreeBSD 11 64-bit	Yes

OS	OS Release	Advanced Monitoring Agent
	FreeBSD 12 64-bit	Yes
	FreeBSD 13 64-bit	Yes

9.5.4 VM Failure Management

A virtual machine failure occurs when a Windows VM experience a blue screen or a Linux VM becomes unresponsive due to excessive load or other issues during operation. You can set an automatic response action for VM failures.

Prerequisites

- You can set a failure response policy for VMs with the x86_64 CPU architecture
- Before setting a failure response policy, install VMTools on the VM and make sure the VMTools is running properly. For more information about VMTools, see [Virtual Machine VMTools](#).
- The failure response policy takes effect immediately after configuration and does not require a VM reboot.

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 failure response policy as needed.
 - No Action (default): Maintains current state without intervention.
 - Reboot: Automatically reboots the VM. Stops after 5 reboot attempts within 30 minutes.
 - Shut Down: Automatically shuts down the VM.



Note:

For Linux VMs, disable the kdump module upon the next boot after enabling this feature.

4. Click **OK**.

9.5.5 VM Time Synchronization

Time is the benchmark for measuring the development stage and status of all things. In the platform, ensuring that VM time is synchronized with the host time is critically important for business operations. NexaSphere offers two time synchronization mechanisms: BIOS Clock Synchronization and VM Time Synchronization. The following sections introduce these two synchronization mechanisms and their setup methods:

- [BIOS Clock Synchronization](#)
- [Time Synchronization](#)

BIOS Clock Synchronization

BIOS clock synchronization is designed for Windows VMs. For Windows VMs, the hardware clock time comes from the host, and the system time is equivalent to the hardware clock time. Therefore, you can enable BIOS clock synchronization to ensure that the hardware and system times of Windows VMs match the host's hardware clock time.

When creating a new VM, you can enable BIOS clock synchronization in **Advanced Settings > General Options**. Once enabled, the VM will periodically synchronize with the host's BIOS clock automatically.

Time Synchronization

The time synchronization feature is designed to synchronize the system time of Windows and Linux VMs with the host system time. For Linux VMs, the hardware clock time comes from the host, while the system time is calculated independently and only synchronized with the hardware clock at specific points (such as when the system starts), remaining independent otherwise.

The VM time synchronization mechanism includes:

- **Network:** Unlike NTP synchronization, this mechanism does not rely on a specific external network; it only involves communication between the host and VM.
- **Time Zone:** The VM periodically synchronizes its time zone with that of the host to maintain consistency.
- **Synchronization Interval:** By default, synchronization occurs every 60 seconds. You can modify the interval via CLI with options including: 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), 86400 seconds (1 day).

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

- **Synchronization Policy:** The current policy ensures immediate consistency. Regardless of whether the VM's time is ahead or behind, once the synchronization mechanism triggers, the VM's time will immediately align with the host's.

To configure time synchronization for an individual VM, follow these steps:

1. Ensure that Qemu Guest Agent (QGA) is installed on the VM and that QGA is running. For more information, see [Virtual Machine VMTools](#).
2. Disable other time source synchronization policies on the VM (recommended).
3. On the VM page, click **Actions > Advanced Settings > Modify VMTools** to enable time synchronization. Once enabled, the VM will automatically synchronize with the host system time every 60 seconds.

9.5.6 VM User Data

NexaSphere supports importing User Data (custom user-defined data). By providing custom parameters or scripts, you can perform customized configurations or accomplish specific tasks on the VMs.

9.5.6.1 Import User Data to Linux VMs

Prerequisites

- Before importing User Data, ensure the VM image has cloud-init installed. Recommended version: 0.7.9, 17.1, 19.4, or later version.
- When setting hostname and SSH password via User Data, avoid configuring these parameters again on the platform to prevent conflicts.
- For Linux VMs created from images with pre-installed cloud-init, you must import User Data. Otherwise, cloud-init task will wait until timeout.

Procedure

1. Prepare the User Data script for your desired functionality.

Example script:

```
#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

```

This example script performs the following actions:

1. Create user "test" with SSH Key authentication during VM creation.
2. Write to `/etc/hosts`, create directory `/tmp/temp`, and generate files with specified content during boot.
3. Set hostname, enable root user, allow SSH password authentication, and change root password.
4. Run the `echo ls -l /` command.
2. In the navigation pane, choose **Inventory**.
3. In the inventory, right-click a cluster, host, or image, then select **New Virtual Machine**.
4. In the **New Virtual Machine** dialog, select **Advanced Settings > General Options**.
5. Paste your script into the **User Data** input field.
6. Complete the VM creation process to import the User Data.

9.5.6.2 Import User Data to Windows VMs

Prerequisites

- Before importing User Data, ensure the VM image has Cloudbase-init installed. Recommended version: 0.9.11. For more information about Cloudbase-init, see [Cloudbase Documentation](#).
- When setting hostname and SSH password via User Data, avoid configuring these parameters again on the platform to prevent conflicts.
- For Windows VMs created from images with pre-installed Cloudbase-init, you must import User Data. Otherwise, Cloudbase-init task will wait until timeout.

Procedure

1. Prepare the User Data script for your desired functionality.

Example script:

```

#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/zMxAgCISQyAgAAAA==
  path: C:\gzip
  permissions: '0644'

```

This example script creates three files during VM startup: *b64*, *b64_1*, *gzip* on the C drive.

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

9.5.7 VM Template

A virtual machine template is a collection of the virtual machine's operating system, applications, and configuration features. Through virtual machine templates, you can quickly deploy multiple virtual machines with the same properties, improving efficiency and consistency.

9.5.7.1 Clone a Virtual Machine to a Template

Clone a virtual machine to a template. The source virtual machine remains available for use after it has been cloned into a template.

Prerequisites

- Make sure the virtual machine is running, paused, shut down, or crashed.
- Detach all shared disks and RDM disks on the virtual machine.

Procedure

1. In the navigation pane, choose **Inventory > VM and Host**.
2. Select the target VM and click **Actions > Clone > Clone to Template**.
3. In the **Clone Virtual Machine to Template** dialog, set the following parameters:
 - **Name:** Enter the template name.



Note:

The name must be unique and cannot duplicate the names of existing virtual machines or templates.

- **Data Storage:** Location where the template will be stored. By default, this is consistent with the source virtual machine's data storage.
- **Tag:** Attach tags to identify different templates. For more information, see [Tag Management](#)

**Note:**

If the VM has GPU, USB, or PCIe devices attached (except for PCIe devices automatically attached by SR-IOV NICs), the cloned template will not include these devices. To keep these devices, you can convert the VM to a template.

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

What's next

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

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

9.5.7.2 Convert a Virtual Machine to a Template

Convert a powered-off virtual machine to a template. After conversion, the source VM becomes template-only and will be removed from the virtual machine navigation tree. The template will appear in the **Image and Template > Template File** navigation tree.

Prerequisites

- Make sure the virtual machine is shut down.
- Detach all shared disks and RDM disks on the virtual machine.
- Make sure the virtual machine has not associated with any backup plans.

Procedure

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

What's next

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

9.5.7.3 Convert a Template to a Virtual Machine

Convert a template directly to a virtual machine. Use this when you need to modify template configurations or no longer require the VM to server as a deployment template.

Procedure

1. In the navigation pane, choose **Inventory > Image and Template**.
2. Select the target template in **Template File** and click **Actions > Convert to Virtual Machine**.
3. In the **Convert to Virtual Machine** dialog, set the following parameters:
 - **Name:** Enter a name for the virtual machine.
 - **Power Status:** Choose whether to automatically power on the VM after conversion.
 - **Location:** Automatically assigned by default. Select a target host or cluster for the VM.
 - For local storage scenario, you cannot specify the VM location. By default, the VM is allocated to the host where local disks reside.
 - For non-local storage scenarios, you can specify the VM location.
4. Review the configuration and click **OK**.

What's next

After converting a template to a virtual machine, you can make adjustments to the virtual machine configuration as needed and then convert the virtual machine back into a template.

9.5.7.4 Create a Virtual Machine from a Template

Create identical virtual machines from a template. You'll be able to customize hardware, software, and other configurations.

Prerequisites

- The platform provides system parameters for VMs to control default settings globally. Before creating a VM, you can customize VM-related settings in the system parameters to control the default VM features. For more information, see [System Parameters](#).
- Before creating a virtual machine based on a template, make sure there is already a virtual machine template available in the platform.

Procedure

1. In the navigation pane, choose **Inventory > VM and Host**.
2. Right-click a valid parent resource (cluster or host) of virtual machines and select **New Virtual Machine**.
3. In the **Select VM Creation Type** dialog, choose **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, set the following parameters:

a) Complete the template information configuration.

- **Template:** Select a template and create a virtual machine based on this template.

b) Complete the basic information configurations.

- **Name:** The name of the virtual machine
- **Quantity:** The number of virtual machines to create this time



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
- c) Complete the hardware information configurations.
- **CPU:** Supports adjusting the total number of cores and the number of cores per socket, and setting CPU hot plug.
 - **Memory:** Supports adjusting the memory size and setting memory resource priority.
 - **Disk:** Supports modifying the cache mode of the disk and setting AIO acceleration. You can add a new hard disk to the virtual machine by clicking **Add Hardware > Disk**. The new hard disk allows customization of its capacity and properties.
 - **NIC:** Supports modifying the NIC model, port group, NIC queue number, MAC address, IP address, and DNS assignment. You can add a new NIC to the virtual machine by

clicking **Add Hardware > NIC**. The new NIC allows customization of its address and properties.

- **CD/DVD Drive**: Supports loading ISO image files onto the virtual machine for booting from an ISO optical drive.
- **GPU Device 1**: Loads a GPU device onto the virtual machine, supporting physical GPU devices and vGPU devices.

You can add a GPU device by clicking **Add Hardware > GPU Device**.

- **USB Device 1**: Loads a USB device onto the virtual machine, supporting direct connection and redirection.

You can add a USB device by clicking **Add Hardware > USB Device**. A single virtual machine supports adding up to 1 USB device.

- **Other Hardware**: Does not support modifying the graphics card and sound card configuration in the template file.

d) Complete the advanced settings.

General Options

- **Description**: Displays the description recorded in the template. You can modify the VM description.
- **Tag**: Displays the tag recorded in the template. You can customize different 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 VM configurations require VMTools. After VM creation, it is recommended to install VMTools to enable certain configurations. For more information about VMTools, see [Virtual Machine VMTools](#).

9.5.7.5 Manage Virtual Machine Templates

You can modify virtual machine templates according to your business needs, including editing names and descriptions, changing configurations, and altering ownership. If you no longer need a template, you can delete the virtual machine template.

Edit Name and Description

If you only need to modify the template's name and description, in the target template page, click on the **Action > Edit Name and Description** to make the changes.

**Note:**

The new name must be unique and cannot duplicate the names of existing virtual machines or templates.

Modify Configuration

If you need to modify the basic information, total number of CPU cores, number of cores per socket, memory size, or hard disk capacity of the template, in the target template page, click on the **Action > Modify Configuration** and make the necessary changes to the corresponding settings.

If you need to make extensive adjustments to the template configuration, you can first convert the template to a virtual machine, then adjust the configuration as needed. Once the adjustments are complete, convert the virtual machine back into a template for future use.

Change Ownership

If you need to change the owner of the template, in the target template page, click on the **Action > Change Ownership** and specify the new owner.

Delete Virtual Machine Template

You can delete a virtual machine template using either of the following methods:

- To delete a single template: Navigate to the target template page, click on the **Action > Delete** to perform the deletion operation.
- To delete multiple templates in bulk: Navigate to the parent resource (data center) corresponding **Image and Template** sub-page, select the target templates, and then click on the **Bulk Action > Delete** to perform the deletion operation.

**Note:**

After deleting a template, you will not be able to recover the template files (including network card, hard disk files, and database records). Proceed with caution.

9.5.8 VM Specification

You can create and manage customization specifications for Windows and Linux VM operating systems. VM specifications are files that contain VM operating system settings.

When you apply a VM specification during batch VM deployments from the template, you can prevent potential conflicts such as duplicated VM hostnames or SIDs.

9.5.8.1 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 details page of the root node, click **VM Specification**.
4. On the **VM Specification** tab, click **New VM Specification**.
5. In the **New VM Specification** dialog, set the following parameters:
 - **Name**: Enter a name for the VM 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.

- 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

- **Hostname Configuration:** 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 a administrator password.

**Note:**

The password must be 8 to 16 characters in length and contain at least two of the following: lowercase letters, uppercase letters, numbers, or special characters.

- **Confirm Password:** Confirm the password by typing it again.

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

9.5.8.2 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 details page of the root node, click **VM Specification**.
4. On the **VM Specification** tab, click **New VM Specification**.
5. In the **New VM Specification** dialog, set the following parameters:
 - **Name:** Enter a name for the VM specification.
 - **Description:** Enter a brief description for the VM specification.
 - **Target VM OS:** Select **Linux**.
 - **Hostname Configuration:** 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.
 - 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 a administrator password.



Note:

The password must be 8 to 16 characters in length and contain at least two of the following: lowercase letters, uppercase letters, numbers, or special characters.

- **Confirm Password:** Confirm the password by typing it again.

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

9.5.8.3 Manage VM Specifications

Procedure

1. In the navigation pane, choose **Inventory > VM and Host**.
2. Select the root node.
3. On the details page of the root node, click **VM Specification**.
4. On the **VM Specification** tab, select a target VM specification.
 - Click **Actions > Edit Name and Description** to modify the name and description of the VM specification.
 - Click **Actions > Modify Configuration** to modify how the virtual machine participates in network, hostname configuration, and administrator password.
 - Click **Actions > Delete** to delete a VM specification.

9.5.9 VM Migration Management

NexaSphere provides both manual migration and HA automatic migration mechanisms. For more information about HA automatic migration, see [VM HA](#). This section introduces the following three scenarios for virtual machine manual migration:

- [Change Host](#): Migrate a virtual machine to another host. This action only changes the host where the virtual machine runs.
- [Change Data Storage](#): Migrate a virtual machine to another data storage. This action only changes data storage.
- [Change Host and Data Storage](#): Migrate a virtual machine to another host and data storage.

9.5.9.1 Change Host

9.5.9.1.1 Hot Migration

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

Prerequisites

- The virtual machine's data storage is local storage, NFS, nSDS distributed storage, SAN storage, or ZHPS distributed storage.
- Before you hot migrate a virtual machine, detach the data disks, ISOs, or peripheral devices (if any) from the virtual machine first. If the virtual machine uses a local storage, you can hot migrate the virtual machine with local data disks attached. If the virtual machine uses a shared storage, you can hot migrate the virtual machine with local data disks, shared disks, and ISOs attached.
- If vNUMA is enabled for a virtual machine, make sure the pNUMA architecture of the destination host is consistent with that of the source host.
- Before you hot migrate a virtual machine with VF NICs attached, make sure:
 - The virtual machine has VMTools installed and the VMTools is running.
 - VF NIC VM migration may fail if the internal driver version of the VM is too low. Check and update the driver version as needed
- If the selected virtual machine has associated with VM scheduling policies, modify these policies based on your business needs to avoid policy conflicts.

Procedure

1. In the navigation pane, choose **Inventory > VM and Host**.
2. Select the target virtual machine and click **Actions > Migration Management > Change Host**.
3. In the **Change Host** dialog, choose a destination host.

If the migration is blocked because the VM has been high-loaded for a long time, you can enable auto-converge to improve the success rate of the migration.

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



Note:

If the virtual machines have VF NICs attached, depending on the actual network environment, short network interruptions may occur during the migration of VMs with VF NICs attached. Proceed with caution.

What's next

After the migration, you can re-attach those data disks, ISOs, or peripheral devices to the virtual machine.

9.5.9.1.2 Cold Migration

Cold migration migrates a powered-off virtual machine.

Prerequisites

- The virtual machine's data storage is local storage.
- Before you cold migrate a virtual machine, detach the data disks, ISOs, or peripheral devices (if any) from the virtual machine first.



Note:

If the virtual machine uses a local storage, you can cold migrate the virtual machine with local data disks attached.

- If the selected virtual machine has associated with VM scheduling policies, modify these policies based on your business needs to avoid policy conflicts.

Procedure

1. In the navigation pane, choose **Inventory > VM and Host**.
2. Select the target virtual machine and click **Shut Down**.
3. After shutdown, click **Actions > Migration Management > Change Host**.
4. In the **Change Host** dialog, choose a destination host.
5. Review the configuration and click **OK**.

What's next

After the migration, you can re-attach those data disks, ISOs, or peripheral devices to the virtual machine.

9.5.9.2 Change Data Storage

9.5.9.2.1 Hot Migration from SAN Storage to SAN Storage

Prerequisites

- Make sure that the source and destination data storage are in the same cluster.
- Before you migrate a virtual machine, detach the ISOs or LUN devices from the virtual machine first.
- If a shared disk is attached to the virtual machine, the shared disk cannot be migrated. Detach the shared disk first.
- Before you hot migrate a virtual machine with VF NICs attached, make sure:

- The virtual machine has VMTools installed and the VMTools is running.
- VF NIC VM migration may fail if the internal driver version of the VM is too low. Check and update the driver version as needed

Procedure

1. In the navigation pane, choose **Inventory > VM and Host**.
2. Select the target virtual machine and click **Actions > Migration Management > Change Data Storage**.
3. In the **Change Data Storage** dialog, choose a destination data storage and set a migration bandwidth as needed.
4. Review the configuration and click **OK**.



Note:

- If the virtual machines have VF NICs attached, depending on the actual network environment, short network interruptions may occur during the migration of VMs with VF NICs attached. Proceed with caution.
- If the virtual machines have snapshots, migrating storage while the VMs are running will delete these snapshots. Proceed with caution. To keep snapshots, shut down VMs before the migration.

What's next

After the migration, you can re-attach those data disks, ISOs, or peripheral devices to the virtual machine.

9.5.9.2.2 Cold Migration from SAN Storage to SAN Storage

Prerequisites

- Make sure that the source and destination data storage are in the same cluster.
- Before you migrate a virtual machine, detach the ISOs or LUN devices from the virtual machine first.
- If a shared disk is attached to the virtual machine, the shared disk cannot be migrated. Detach the shared disk first.

Procedure

1. In the navigation pane, choose **Inventory > VM and Host**.
2. Select the target virtual machine and click **Shut Down**.

3. After shutdown, click **Actions > Migration Management > Change Data Storage**.
4. In the **Change Data Storage** dialog, choose a destination data storage.
5. Review the configuration and click **OK**.

What's next

After the migration, you can re-attach those data disks, ISOs, or peripheral devices to the virtual machine.

9.5.9.2.3 Hot Migration Between SAN Storage and nSDS Distributed Storage

Prerequisites

- Make sure that the source and destination data storage are in the same cluster.
- Before you migrate a virtual machine, detach the ISOs or LUN devices from the virtual machine first.
- If a shared disk is attached to the virtual machine, the shared disk cannot be migrated. Detach the shared disk first.
- Before you hot migrate a virtual machine with VF NICs attached, make sure:
 - The virtual machine has VMTTools installed and the VMTTools is running.
 - VF NIC VM migration may fail if the internal driver version of the VM is too low. Check and update the driver version as needed

Procedure

1. In the navigation pane, choose **Inventory > VM and Host**.
2. Select the target virtual machine and click **Actions > Migration Management > Change Data Storage**.
3. In the **Change Data Storage** dialog, choose a destination data storage.
If you hot migrate a virtual machine from a SAN storage to a nSDS distributed storage, you can specify a root disk pool and data disk pool for the disks to be migrated.
4. Review the configuration and click **OK**.



Note:

- If the virtual machines have VF NICs attached, depending on the actual network environment, short network interruptions may occur during the migration of VMs with VF NICs attached. Proceed with caution.

- If the virtual machines have snapshots, migrating storage while the VMs are running will delete these snapshots. Proceed with caution. To keep snapshots, shut down VMs before the migration.

What's next

After the migration, you can re-attach those data disks, ISOs, or peripheral devices to the virtual machine.

9.5.9.3 Change Host and Data Storage

9.5.9.3.1 Hot Migration Across Data Storage of the Same Type

Prerequisites

- Make sure the data storage where the virtual machine resides meet the storage combination requirements.

Supported storage combination:

- Hot migration from local storage to local storage
- Hot migration from nSDS distributed storage to nSDS distributed storage
- Hot migration from NFS to NFS
- Hot migration from SAN storage to SAN storage
- Before you migrate a virtual machine, detach the shared disks, ISOs, or peripheral devices (if any) from the virtual machine first.
- If vNUMA is enabled for a virtual machine, make sure the pNUMA architecture of the destination host is consistent with that of the source host.
- Before you hot migrate a virtual machine with VF NICs attached, make sure:
 - The virtual machine has VMTTools installed and the VMTTools is running.
 - VF NIC VM migration may fail if the internal driver version of the VM is too low. Check and update the driver version as needed
- If the selected virtual machine has associated with VM scheduling policies, modify these policies based on your business needs to avoid policy conflicts.
- For hot migration across data storage of the same type, you can only migrate the entire virtual machine.
- Before you hot migrate a virtual machine across distributed storage, make sure that the monitor nodes of these two distributed storage can communicate with each other.

Procedure

1. In the navigation pane, choose **Inventory > VM and Host**.
2. Select the target virtual machine 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 set a migration bandwidth as needed.
 - For hot migration from nSDS distributed storage to nSDS distributed storage, you can specify disk 1 storage pool and choose whether to migrate the attached data disks.
 - If the migration is blocked because the VM has been high-loaded for a long time, you can enable auto-converge to improve the success rate of the migration.
4. Review the configuration and click **OK**.

**Note:**

- If the virtual machines have VF NICs attached, depending on the actual network environment, short network interruptions may occur during the migration of VMs with VF NICs attached. Proceed with caution.
- If the virtual machines have snapshots, migrating storage while the VMs are running will delete these snapshots. Proceed with caution. To keep snapshots, shut down VMs before the migration.

9.5.9.3.2 Hot Migration Across Data Storage of Different Types

Prerequisites

- Make sure the data storage where the virtual machine resides meet the storage combination requirements.

Supported storage combination:

- Hot migration between nSDS distributed storage and SAN storage
- Hot migration between nSDS distributed storage and NFS storage
- Hot migration between local storage and SAN storage
- Hot migration between local storage and nSDS distributed storage
- Hot migration between local storage and NFS storage
- Hot migration between SAN storage and NFS storage
- Before you migrate a virtual machine, detach the shared disks, ISOs, or peripheral devices (if any) from the virtual machine first.

- If vNUMA is enabled for a virtual machine, make sure the pNUMA architecture of the destination host is consistent with that of the source host.
- Before you hot migrate a virtual machine with VF NICs attached, make sure:
 - The virtual machine has VMTools installed and the VMTools is running.
 - VF NIC VM migration may fail if the internal driver version of the VM is too low. Check and update the driver version as needed
- If the selected virtual machine has associated with VM scheduling policies, modify these policies based on your business needs to avoid policy conflicts.
- For hot migration across data storage of different types, you can only migrate the entire virtual machine.

Procedure

1. In the navigation pane, choose **Inventory > VM and Host**.
2. Select the target virtual machine 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 set a migration bandwidth as needed.
 - If you hot migrate a virtual machine from a SAN storage, local storage, or an NFS storage to a nSDS distributed storage, you can specify a disk 1 pool for the disks to be migrated.
 - If the migration is blocked because the VM has been high-loaded for a long time, you can enable auto-converge to improve the success rate of the migration.
4. Review the configuration and click **OK**.



Note:

- If the virtual machines have VF NICs attached, depending on the actual network environment, short network interruptions may occur during the migration of VMs with VF NICs attached. Proceed with caution.
- If the virtual machines have snapshots, migrating storage while the VMs are running will delete these snapshots. Proceed with caution. To keep snapshots, shut down VMs before the migration.

9.5.9.3.3 Hot Migration Across Storage Pools Within the Same nSDS Distributed Storage

Prerequisites

- Make sure the data storage where the virtual machine resides is distributed storage.
- Before you migrate a virtual machine, detach the shared disks, ISOs, or peripheral devices (if any) from the virtual machine first.
- If vNUMA is enabled for a virtual machine, make sure the pNUMA architecture of the destination host is consistent with that of the source host.
- Before you hot migrate a virtual machine with VF NICs attached, make sure:
 - The virtual machine has VMTTools installed and the VMTTools is running.
 - VF NIC VM migration may fail if the internal driver version of the VM is too low. Check and update the driver version as needed

Procedure

1. In the navigation pane, choose **Inventory > VM and Host**.
2. Select the target virtual machine 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 set a migration bandwidth as needed.
 - For hot migration across storage pools within the same nSDS distributed storage, you can specify the disk 1 storage pool and choose whether to migrate the attached data disks.
 - If the migration is blocked because the VM has been high-loaded for a long time, you can enable auto-converge to improve the success rate of the migration.
4. Review the configuration and click **OK**.



Note:

- If the virtual machines have VF NICs attached, depending on the actual network environment, short network interruptions may occur during the migration of VMs with VF NICs attached. Proceed with caution.
- If the virtual machines have snapshots, migrating storage while the VMs are running will delete these snapshots. Proceed with caution. To keep snapshots, shut down VMs before the migration.

9.5.9.3.4 Clod Migration Across Data Storage of the Same Type

Prerequisites

- Make sure the data storage where the virtual machine resides meet the storage combination requirements.

Supported storage combination:

- Cold migration between nSDS distributed storage and nSDS distributed storage
- Cold migration between NFS storage and NFS storage
- Detach all data disks, shared disks, ISOs, or peripheral devices (if any) from the virtual machine before the migration.
- If the selected virtual machine has associated with VM scheduling policies, modify these policies based on your business needs to avoid policy conflicts.
- For cold migration between nSDS distributed storage and nSDS distributed storage: If you set a cold migration network for the target data storage, the virtual machine is migrated through the cold migration network. Otherwise, the virtual machine is migrated through the management network.
- For cold migration between NFS storage and NFS storage: The destination NFS storage can be attached to the cluster where the virtual machine to be migrated is located.

Procedure

1. In the navigation pane, choose **Inventory > VM and Host**.
2. Select the target virtual machine and click **Shut Down**.
3. After shutdown, 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.
5. Review the configuration and click **OK**.

9.5.9.3.5 Cold Migration Across Storage Pools Within the Same nSDS Distributed Storage

Prerequisites

Detach all data disks, shared disks, ISOs, or peripheral devices (if any) from the virtual machine before the migration.

Procedure

1. In the navigation pane, choose **Inventory > VM and Host**.
2. Select the target virtual machine and click **Shut Down**.
3. After shutdown, 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.
5. Review the configuration and click **OK**.

9.5.10 VM Performance

If your business requires VMs to have certain characteristics, such as a stronger ability to compete for computing resources, you can achieve this by configuring the relevant VM settings provided by NexaSphere.

You can set these configurations when creating a new VM or modify the settings of existing VMs. This section assumes that you have already created several VMs.

The relevant VM configurations include:

- [VM Resource Contention](#)
- [VM CPU](#)
- [VM QoS](#)

9.5.10.1 VM Resource Contention

NexaSphere provides the following three configurations to enhance a virtual machine's resource contention capability:

- [CPU NUMA Binding Configuration](#)
- [EmulatorPin Configuration](#)
- [Configure CPU and Memory Resource Priorities](#)

CPU NUMA Binding Configuration

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

If you encounter either of the following two business scenarios, consider using the CPU binding feature to enhance VM performance:

- CPU-intensive small application scenario:

CPU-intensive applications have a high demand for CPU resources. If numerous such applications run on VMs, it can lead to competition for CPU resources. CPU binding ensures each application runs on specific physical CPUs, thus avoiding resource contention and improving system performance.

- Uneven pressure on multi-core CPU scenario:

For cases where multiple applications concentrate on one or a few CPUs, you can manually adjust the load on each CPU through CPU binding, with changes taking effect immediately.

To configure CPU binding, follow these steps:

1. On the target VM page, click **Shut Down** to power off the VM.
2. Click **Modify Configuration**, then select **Hardware Information > CPU > CPU NUMA Binding**, choose between automatic intelligent binding or manual binding as needed:

- **Automatic Intelligent Binding:** Binds the VM's vCPUs to pCPUs within pNUMA nodes in descending order of pNUMA node ID. When all pCPUs in a pNUMA node are bound by the VM's vCPUs, it will proceed to bind vCPUs to pCPUs in the next pNUMA node. If all pCPUs are bound but there are still unbound vCPUs, it will cycle back to the first bound pNUMA node to continue binding.
- **Manual Binding:** Manually bind vCPUs to pCPUs according to the host's pNUMA topology, ensuring all vCPUs are bound.
 - A single vCPU can be bound to one or more pCPUs, and a single pCPU can be bound by one or more vCPUs, provided that all pCPUs bound to a single vCPU reside within the same pNUMA node.
 - Displays the average usage rate of each pCPU over the past 15 minutes to assist in selecting the optimal pCPU for binding.



Note:

- If you have set up CPU overcommitment, the number of vCPUs may exceed the number of pCPUs. In this case, it is recommended not to have more vCPUs than bound pCPUs, as this could significantly impact VM performance.
- After binding the CPU NUMA, if you change the number of virtual machine CPU cores, the CPU NUMA binding will be automatically canceled.

3. After clicking **OK**, the CPU NUMA binding will be completed. You can then click **Power On** to restart the VM.

EmulatorPin Configuration

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.

If you encounter the following business scenario, consider using EmulatorPin:

- Parallel multi-business scenario:

In a scenario where multiple VMs run on a single host with different services on each VM, this can lead to varying degrees of resource consumption. EmulatorPin achieves isolation of primary service processes across different VMs by binding the QEMU main thread of each VM to specific pCPUs, ensuring relatively stable system performance.

To configure EmulatorPin binding, follow these steps:

1. On the target VM page, click **Actions > Advanced Settings > Modify Other Options**, and select **EmulatorPin**: bind the threads other than vCPU and IO threads in the VM to the host's pCPUs according to the host's pNUMA structure.
2. After clicking **OK**, the EmulatorPin binding will be completed.

Configure CPU and Memory Resource Priorities

If Virtual Machine A shares the CPU and memory resources of Host A with other virtual machines, and its business importance and priority are higher than those of other virtual machines. In scenarios where resources are tight, such as when the host's load rate is too high, these virtual machines may compete for host resources. In this case, you can set the CPU resource priority and memory resource priority of Virtual Machine A to **High**, while others remain at **Normal**, to enhance Virtual Machine A's capability to acquire resources.

To configure resource priority, follow these steps:

1. In the target virtual machine page, click **Modify Configuration**, then choose **Hardware Information > CPU > CPU Resource Priority**.
2. Click **OK** to complete configuration.

9.5.10.2 VM CPU

NexaSphere provides CPU mode configuration functionality, which you can use to set the virtual machine's CPU model to match the host's CPU model. This allows the virtual machine to inherit some or all of the host's CPU features, meeting specific business requirements.

- [Introduction to Virtual Machine CPU Modes](#)
- [Configuring Virtual Machine CPU Mode](#)

Introduction to Virtual Machine CPU Modes

NexaSphere supports the following CPU mode settings:

- **None (Default):** The virtual machine's CPU model is emulated by QEMU, inheriting a limited set of features from the host CPU. This mode is recommended for migration scenarios.
- **Compatible:** The virtual machine's CPU model closely matches or is identical to the host CPU model, for example, both showing as Haswell Intel CPUs. Compared to the None mode, this mode allows the virtual machine to inherit more features from the host CPU and can be used in migration scenarios. Note that this setting is not supported on aarch architecture clusters and virtual machines.
- **Passthrough:** The virtual machine's CPU model is identical to the host CPU model. Additionally, the virtual machine inherits all CPU features from the host, such as extended page tables, large memory pages, and virtualization support. Compared to None, Compatibility, and Custom modes, this mode inherits the most features from the host CPU and is suitable for business scenarios with higher requirements for virtual machine functionality.
- **Custom (a specific CPU model):** The virtual machine is configured with the specified custom CPU model. After configuring a custom CPU model, the virtual machine may have different CPU features compared to its previous configuration.

Configuring Virtual Machine CPU Mode

NexaSphere offers configuration settings at both the virtual machine level and the cluster level. The precedence order for these settings is: virtual machine level > cluster level. This section describes how to configure the CPU mode for a single virtual machine:

1. In the target virtual machine page, click the **Shut Down** button to power off the virtual machine.
2. Click **Modify Configuration**, then choose **Hardware Information > CPU > CPU Mode**. Select the desired CPU mode according to your needs.
3. Click **OK** to complete the CPU mode configuration. You can then click the **Power On** button to restart the virtual machine.

9.5.10.3 VM QoS

QoS (Quality of Service) resolves issues such as network latency and congestion by setting IO bandwidth thresholds. When the network is overloaded or congested, QoS ensures that traffic is not delayed or dropped, maintaining efficient network operation.

You can set QoS to limit disk bandwidth and inbound and outbound network bandwidth, ensuring that IO bandwidth does not exceed the set threshold. If you do not configure QoS, there will be no restrictions on IO bandwidth. This section primarily describes how to set up disk QoS and network QoS:

- [Configure Disk QoS](#)
- [Configure Network QoS](#)

Configure Disk QoS

To set up disk QoS, follow these steps:

1. In the target virtual machine page, click **Modify Configuration**, then choose **Hardware Information > Disk {\$n} > QoS**. Turn on the QoS switch to configure QoS for disk n (for example, Disk 1).

Disk QoS supports configuring bandwidth limits and IOPS limits. After turning on the switch, you must configure at least one of these limits:

- **Bandwidth Limit:** Set the upper limit of read/write speed per second for the disk. Basic units include MB/s and GB/s. It is recommended not to set this value too low to avoid abnormal operation of the virtual machine.
 - **Total Speed:** Set the upper limit of total read/write speed for the disk.
 - **Read/Write Speed:** Set separate upper limits for read speed and write speed.
- **IOPS Limit:** Set the upper limit of read/write operations per second for the disk.
 - **Total IOPS:** Set the upper limit of total read/write IOPS for the disk.
 - **Read/Write IOPS:** Set separate upper limits for read IOPS and write IOPS.

2. Click **OK** to complete the disk QoS configuration.

Configure Network QoS

To configure network QoS, follow these steps:

1. In the target virtual machine page, click **Modify Configuration**, then choose **Hardware Information > NIC {\$n} > QoS**. Turn on the QoS switch to configure QoS for NIC n (for example, NIC 1).

QoS supports configuring outbound and inbound bandwidth limits. Basic units include Kbps, Mbps, and Gbps. After turning on the switch, you must configure at least one of these bandwidth limits:

- **Transmit Bandwidth:** The upper limit of upload bandwidth from the virtual machine.
- **Receive Bandwidth:** The upper limit of download bandwidth to the virtual machine.

2. Click **OK** to complete the network QoS configuration.

9.5.11 VM Snapshot Management

9.5.11.1 Overview

Snapshot: A snapshot is a point-in-time capture of data status in a disk. Before performing business-sensitive operations, you can create a snapshot at specified time points to record the state of the system disk and data disks of a virtual machine. This allows for quick rollback in case of breakdowns.

Snapshot Types

- **Manual Snapshot:** Users can manually create a snapshot of the entire virtual machine at any time. For more information, see [Snapshot Basic Operations](#).
- **Scheduled Snapshot:** The system automatically executes snapshots at specified times according to a snapshot policy. For more information, see [Snapshot Policy Basic Operations](#).
- **Automatic Snapshot:** The system triggers one-time automatic snapshots in specific scenarios, including system reset and virtual machine cloning. When creating an image, a temporary snapshot is automatically created and will be automatically deleted once the image creation is complete.

Scenarios

- **Fast recovery from breakdowns:** If the production environment breaks down, you can rollback snapshots to return to the normal state. Snapshot rollback is a temporary use for unexpected breakdowns. To backup data for a long term, we recommend that you use the Backup Service.
- **Data exploration:** You can create snapshots for production data and use the snapshots to do data mining, query, and development and test.
- **Fault tolerance improvement:** Before you perform a business-sensitive operation such as upgrading your system or migrating data, we recommend that you create one or more snapshots. This allows fast recovery to the normal state in case of any errors or exceptions during the operation.

9.5.11.2 Snapshot Basic Operations

9.5.11.2.1 Create a Snapshot

You can create one or more snapshots for a virtual machine to retain its temporary state at a specific point in time, allowing for quick rollback in case of a fault.

Prerequisites

- To create a snapshot, make sure the virtual machine does not have any shared disks attached.

- You cannot take memory snapshots for virtual machines using distributed storage.
- You cannot take memory snapshots for virtual machines in the Compatible CPU mode.
- To create a memory snapshot, make sure the virtual machine is in the running state and any attached peripheral devices are detached from the virtual machine.
- Excessive snapshots will lower the VM performance, increase data security risks, and occupy data storage space. For long-term data backup, you can use the backup service.

Procedure

1. In the navigation pane, choose **Inventory > VM and Host**.
2. Select a target VM and click **Actions > Snapshot and Image > Create Snapshot**.
3. In the **Create Snapshot** dialog, set the following parameters:
 - **Name:** Enter a name for the snapshot.
 - **Description:** Enter a brief description for the snapshot.
 - **Memory Snapshot:** A memory snapshot captures the real-time state of a virtual machine.
4. Review the configuration and click **OK**.



Note:

To ensure memory consistency, the virtual machine will be briefly paused during memory snapshot creation.

9.5.11.2.2 Revert to a Snapshot

Reverting a VM to the selected snapshot overwrites the existing data.

Procedure

1. In the navigation pane, choose **Inventory > VM and Host**.
2. Select a target VM and click **Snapshot**.
3. On the **Snapshot** tab, select a snapshot point and click **Revert**.
4. In the **Revert Snapshot?** dialog, set the following parameters:
 - **Recovery Method:** Support full recovery and custom recovery
 - Full Recovery: A full recovery restores both the VM data and disk sequence, which may attach the previously detached disks or detach existing disks.
 - Custom Recovery: Select disks for recovery. During recovery, only data from the selected disks is restored, but the disk sequence is not restored.

**Note:**

If a disk in the snapshot has been detached, it is currently not supported for recovery.

- **Power Status:** Choose whether to automatically power on the virtual machine after reverting the snapshot.

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

**Note:**

- Reverting a VM to the selected snapshot automatically powers off the VM and overwrites the existing data.
- If the number of NICs or CD/DVD drives differs from the snapshot configuration, any removed NICs or CD/DVD drives after the snapshot creation will be attached, and any newly added ones after the snapshot creation will be removed.
- When restoring a memory snapshot, if the memory snapshot contains IP or MAC addresses that are currently in use by other VMs, resolve conflicts before restoring the snapshot.

If you force revert, the platform will handle conflicts as follows:

- Conflicting IP addresses: Keeps the IP addresses from the snapshot and you need to manually modify these IP addresses after restoration.
- Conflicting MAC addresses: Keeps the MAC addresses from the snapshot but disables affected NICs after restoration.

9.5.11.2.3 New Virtual Machine from Snapshot

You can create a new virtual machine directly from a snapshot. This recovery method does not overwrite the existing virtual machine.

Prerequisites

- Make sure the virtual machine has at least one successfully generated snapshot.
- Make sure the platform has sufficient compute, storage, and network resources to support the new virtual machine.
- You cannot create a new virtual machine if the selected snapshot is a memory snapshot.

Procedure

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

2. Select a target VM and click **Snapshot**.
3. On the **Snapshot** tab, select a snapshot point and click **New Virtual Machine**.
4. In the **New Virtual Machine from Snapshot** dialog, set the following parameters:

Snapshot Information

- **Snapshot:** Display the selected snapshot.

Basic Information

- **Name:** Name of the virtual machine.
- **Quantity:** Default is 1, modification is not supported.
- **Location:** Host or cluster location where the virtual machine resides.
- **OS:** Operating system of the virtual machine, including Linux and Windows
- **Power Status:** Choose whether to automatically power on the virtual machine after the creation.

Hardware Information

- **CPU:** Support adjusting the total number of cores.
- **Memory:** Support adjusting the memory size.
- **Disk:** Display the disk configuration recorded in the snapshot. Modification is not supported.
- **NIC:** Support adjusting port groups, MAC address, IP address, DNS assignment, and security groups

You can add a new NIC to the virtual machine by clicking **Add NIC**. The new NIC allows customization of the network address and features.

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

What's next

Some VM configurations require VMTools. After VM creation, it is recommended to install VMTools to enable certain configurations. For more information about VMTools, see [Virtual Machine VMTools](#).

9.5.11.2.4 View Snapshot

The platform provides a unified entry point for users to centrally manage all snapshots. Additionally, you can view snapshots of a specific virtual machine on its **Snapshot** tab.

Procedure

1. In the navigation pane, choose **Data Protection > Snapshot**.

2. The **Snapshot** page displays snapshot information in a tree-like hierarchical manner and supports sorting by the number of snapshots or by snapshot size.
 - After selecting a specific virtual machine, the details area shows all snapshots under that VM, including snapshot name, capacity, creation time, and available actions.
 - After selecting a specific snapshot under a VM, the details area displays available actions for that snapshot, basic information, and snapshot details.

9.5.11.2.5 Delete a Snapshot

Procedure

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



Note:

- Deleting the current snapshot also deletes snapshots on its branch.
- Deleting the current snapshot consumes I/O bandwidth. Do not restart the management node service at this time.
- The snapshot data is committed before it is deleted. This process occupies a certain amount of the data storage. If the data storage capacity is insufficient, you might fail to delete the snapshot.

9.5.11.3 Snapshot Policy Basic Operations

9.5.11.3.1 New Snapshot Policy

Associate a snapshot policy to a specified virtual machine to achieve periodic automatic snapshots.

Prerequisites

- A virtual machine can have only one snapshot policy associated. A virtual machine that already has an associated snapshot policy cannot be associated again.
- If a virtual machine has shared disks or RDM disks attached, it cannot be associated with a snapshot policy.
- Make sure the virtual machine uses nSDS distributed storage.

Procedure

1. In the navigation pane, choose **Data Protection > Snapshot Policy**.
2. On the **Snapshot** page, click **New Snapshot Policy**.
3. In the **New Snapshot Policy** dialog, set the following parameters:
 - **Name**: Name of the snapshot policy.
 - **Description**: Optional.
 - **Scheduled Snapshot Cycle**: Set the frequency of snapshot generation. Snapshots will be created according to the specified execution time, including by week and by month. It supports setting more granular snapshot creation times, precise to the minute level.
 - **Start Time**: Set the time to start executing the snapshot policy
 - **End Time**: Set the time to stop snapshot policy, including Never and Custom.
 - **Retained Snapshots**: Set the upper limit for the number of snapshots to be retained, and any excess will be automatically deleted.
 - **Associated VM**: After associating virtual machines, when the snapshot policy takes effect, the system will automatically create snapshots of the associated virtual machines at the specified time.
4. Review the configuration and click **OK**.

9.5.11.3.2 Enable/Disable Snapshot Policy

You can flexibly manage the snapshot policy status, including enabling and disabling. After disabled, the system will pause periodic automatic snapshots until you re-enable the policy.

Procedure

1. In the navigation pane, choose **Data Protection > Snapshot Policy**.
2. On the **Snapshot Policy** page, select the target snapshot and click **Actions > Enable/Disable**.

9.5.11.3.3 Modify Configuration

You can modify the snapshot policy, such as its scheduled snapshot cycle, start and end time, retained snapshots and associated virtual machines.

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, make changes as needed.
4. Review your changes and click **OK**.

9.5.11.3.4 Delete 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**.

9.5.11.4 Snapshots Usage Recommendations

This chapter mainly provides recommendations for using snapshot features in production environments.

1. In a production environment, it is recommended to keep the number of snapshots per disk to no more than five. Too many snapshots can affect the I/O performance, data security, and data storage capacity of the virtual machine/disk.
2. In a production environment, to ensure data integrity, it is not recommended to create snapshots for virtual machines with high I/O. When the virtual machine performs high I/O operations internally, creating a snapshot for the virtual machine means that some data in memory has not yet been written to the disk, and this part of the data will not be saved in the snapshot.
3. Explanation of the storage capacity occupied by snapshots:
 - In local storage or centralized storage scenarios, creating incremental snapshots occupies minimal storage space, while creating full snapshots occupies double the storage space.
 - In distributed storage scenarios, creating a snapshot does not occupy additional storage space by itself, but after creating a snapshot, writing operations to the original disk may trigger Copy-On-Write (COW) for the snapshot, resulting in each snapshot consuming the same amount of storage space as the original disk.
4. Before performing a snapshot rollback operation, it is strongly recommended to create a snapshot for the disk to protect the current state of the disk data.
5. Explanation of the impact of creating a snapshot on current operations:

Generally, there is no impact, but when creating a full snapshot, there may be additional network I/O generated on the storage network, especially in local storage or centralized storage scenarios, where the disk I/O bandwidth is significantly affected, but disk IOPS are largely unaffected.

6. Explanation of the impact of deleting a snapshot on current operations:

Generally, deleting a snapshot also deletes snapshots on its child branches and merges the data into the disk, which can generate additional disk I/O bandwidth, causing the business I/O to potentially slow down slightly.

9.5.12 VM Scheduling Policy

NexaSphere virtual machine scheduling policies enable cross-cluster VM scheduling within the data center. This section introduces the scheduling policy feature and its usage methods:

- [Overview](#)
- [Create Mutually Exclusive/Affinitive VM Scheduling Policies](#)
- [Create VM Mutually Exclusive/Affinitive Host Scheduling Policies](#)
- [Manage VM Scheduling Policies and Related Resources](#)

9.5.12.1 Overview

Virtual Machine Scheduling Policy (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.

Related Definitions

The virtual machine scheduling policy includes the following core components:

- Scheduling Policy: NexaSphere provides four types of scheduling policies, each supporting two execution mechanisms:
 - Policy Type: Supports four types of scheduling policies: VM Exclusive from Each Other, VM Affinitive to Each Other, VMs Exclusive from Hosts, and VMs Affinitive to Hosts. By combining different execution mechanisms, virtual machine scheduling is achieved:
 - VM Exclusive from Each Other: Virtual machines in the same VM scheduling group should not/must not run on the same host.
 - VM Affinitive to Each Other: Virtual machines in the same VM scheduling group should/must run on the same host.
 - VMs Exclusive from Hosts: Given any one of the virtual machines in a VM scheduling group and any one of the hosts in a host scheduling group, the virtual machine should not/must not run the host.

- VMs Affinitive to Hosts: Given any one of the virtual machines in a VM scheduling group and any one of the hosts in a host scheduling group, the virtual machine should/must run the host.
- Execution Mechanism: Supports two mechanisms: Hard and Soft:
 - Hard: Virtual machines are forcibly assigned hosts based on the associated VM scheduling policies. For example, if you associate the VM Exclusive from Each Other policy with a VM scheduling group and select the Hard mechanism for the policy, any two of the virtual machines in the scheduling group are not allowed to run on the same host. If no host is available to be scheduled based on the policy for a virtual machine, the virtual machine will end up failure upon startup.
 - Soft: Virtual machines are primarily assigned hosts based on the associated VM scheduling policies. For example, if you associate the VM Exclusive from Each Other policy with a VM scheduling group and select the Soft mechanism for the policy, any two of the virtual machines in the scheduling group will primarily not run on the same host. If no host is available to be scheduled based on the policy for a virtual machine, the virtual machine will attempt to run on a host that does not satisfy the policy.
- VM Scheduling Group: The basic unit for binding virtual machines to scheduling policies.
 - A virtual machine can only join one virtual machine scheduling group. Once joined, the virtual machine will be assigned to a host according to the scheduling policy associated with the scheduling group.
 - A virtual machine scheduling group can be bound to one or more scheduling policies.
 - A scheduling policy can be bound to one virtual machine scheduling group.
 - Deleting a virtual machine scheduling group will also delete the associated virtual machine scheduling policies.
- Host Scheduling Group: The basic unit for binding hosts to scheduling policies, used only when selecting the virtual machine mutually exclusive hosts or virtual machine host affinity scheduling policies.
 - A host can only join one host scheduling group. Once joined, the host will be scheduled according to the scheduling policy associated with the scheduling group.
 - A host scheduling group can be bound to one or more scheduling policies.
 - A scheduling policy can be bound to one host scheduling group.
 - Deleting a host scheduling group will also delete the associated virtual machine scheduling policies.

Function Principles

NexaSphere supports adding a virtual machine to a virtual machine scheduling group and binding scheduling policies to the group to achieve virtual machine scheduling.


- If a mutually exclusive virtual machine or aggregated virtual machine policy is bound, no host scheduling group needs to be specified, and the virtual machine is assigned to a host according to the policy and its execution mechanism.
- If a virtual machine host affinity or virtual machine mutually exclusive host scheduling policy is bound, the corresponding host scheduling group must be specified, and the virtual machine is assigned to a host according to the policy and its execution mechanism.

The following explains the working principles of the four types of scheduling policies through four scenarios:

Scenario 1: Assume there are three hosts Host1, Host2, and Host3 in the data center. Virtual machine scheduling group A is bound to the **VM Exclusive from Each Other** scheduling policy, and virtual machines VM1 and VM2 have joined this scheduling group and are running on hosts Host1 and Host2, respectively. At this point, virtual machine VM3 joins this scheduling group.


Under different execution mechanisms, the behavior of virtual machine VM3 is as follows:

- Under the mandatory mechanism, virtual machine VM3 adheres to the principle of mandatory mutual exclusion with other virtual machines in the group:
 - If Host3 has sufficient resources, it can normally start and run on Host3.
 - If Host3 does not have sufficient resources, it cannot start on Host3.
- Under the preferred mechanism, virtual machine VM3 adheres to the principle of trying to mutually exclude other virtual machines in the group, prioritizing starting on Host3:
 - If Host3 has sufficient resources, it can normally start and run on Host3.
 - If Host3 does not have sufficient resources, VM3 can attempt to start on another host with sufficient resources. In this scenario, VM3 starts and runs on Host2.

Figure 9-2: VM Exclusive from Each Other (Hard/Soft)

Scenario 2: Assume there are two hosts Host1 and Host2 in the data center. Virtual machine scheduling group A is bound to the **VM Affinitive to Each Other** scheduling policy, and virtual machines VM1 and VM2 have joined this scheduling group and are running on host Host1. At this point, virtual machine VM3 joins this scheduling group. Under different execution mechanisms, the behavior of virtual machine VM3 is as follows:

- Under the mandatory mechanism, virtual machine VM3 adheres to the principle of mandatory aggregation with other virtual machines in the group:
 - If Host1 has sufficient resources, it can normally start and run on Host1.
 - If Host1 does not have sufficient resources, it cannot start on Host1.
- Under the preferred mechanism, virtual machine VM3 adheres to the principle of trying to aggregate with other virtual machines in the group, prioritizing starting on Host1:
 - If Host1 has sufficient resources, it can normally start and run on Host1.
 - If Host1 does not have sufficient resources, VM3 can attempt to start on another host with sufficient resources. In this scenario, VM3 starts and runs on Host2.


Figure 9-3: VM Affinitive to Each Other (Hard/Soft)

Scenario 3: Assume there are three hosts Host1, Host2, and Host3 in the data center. Virtual machine scheduling group A is bound to the **VMs Exclusive from Hosts** scheduling policy, and virtual machines VM1 and VM2 have joined this scheduling group and are running on host Host1. Host scheduling group A is also bound to the **VMs Exclusive from Hosts** scheduling policy, and hosts Host2 and Host3 have joined this scheduling group, each running two virtual machines. At this point, virtual machine VM3 joins virtual machine scheduling group A. Under different execution mechanisms, the behavior of virtual machine VM3 is as follows:

- Under the mandatory mechanism, virtual machine VM3 adheres to the principle of mandatory mutual exclusion with hosts in host scheduling group A:
 - If Host1 has sufficient resources, it can normally start and run on Host1.
 - If Host1 does not have sufficient resources, it cannot start on Host1.
- Under the preferred mechanism, virtual machine VM3 adheres to the principle of trying to mutually exclude hosts in host scheduling group A, prioritizing starting on Host1:
 - If Host1 has sufficient resources, it can normally start and run on Host1.

- If Host1 does not have sufficient resources, VM3 can attempt to start on another host with sufficient resources. In this scenario, VM3 starts and runs on Host2.

Figure 9-4: VMs Exclusive from Hosts (Hard/Soft)




Scenario 4: Assume there are three hosts Host1, Host2, and Host3 in the data center. Virtual machine scheduling group A is bound to the **VMs Affinitive to Hosts** scheduling policy, and virtual machines VM1 through VM5 have joined this scheduling group and are running on hosts Host2 and Host3. Host scheduling group A is also bound to the **VMs Affinitive to Hosts** scheduling policy, and hosts Host2 and Host3 have joined this scheduling group. At this point, virtual machine VM6 joins virtual machine scheduling group A. Under different execution mechanisms, the behavior of virtual machine VM6 is as follows:

- Under the mandatory mechanism, virtual machine VM6 adheres to the principle of mandatory aggregation with hosts in host scheduling group A:
 - If Host2 or Host3 has sufficient resources, it can normally start and run on Host2 or Host3.
 - If Host2 and Host3 do not have sufficient resources, it cannot start on Host2 or Host3.
- Under the preferred mechanism, virtual machine VM6 adheres to the principle of trying to aggregate with hosts in host scheduling group A, prioritizing starting on Host2 or Host3:
 - If Host2 or Host3 has sufficient resources, it can normally start and run on Host2 or Host3.

- If Host2 and Host3 do not have sufficient resources, VM6 can attempt to start on another host with sufficient resources. In this scenario, VM6 starts and runs on Host1.

Figure 9-5: VMs Affinitive to Hosts (Hard/Soft)



Use Cases

You can use the VM Exclusive from Each Other (Hard/Soft) and VMs Affinitive to Hosts (Hard/Soft) functions in the following scenarios.

- Example use case for the VM Exclusive from Each Other (Hard) policy:

Two virtual machines hosting primary and backup databases are required to be deployed on different hosts to ensure high availability of the service.

- For example: Users deploy two virtual machines to host primary and backup MySQL databases, requiring that both virtual machines must be deployed on different hosts to reduce the risk of service downtime. Since deployment is automated, users cannot predict in advance which hosts have available resources. Using the mutually exclusive virtual machine (mandatory) policy, two different hosts can be selected to run these two virtual machines, ensuring high availability of the service.
- Example use case for the VM Exclusive from Each Other (Soft) policy:

Hadoop nodes with different roles are preferably deployed on different hosts to improve overall system performance.

- For example: Users deploy a Hadoop system, for which nodes with different roles such as namenode, datanode, jobtracker, and tasktracker, cannot be predicted in advance regarding the total number of nodes. However, deploying them on different hosts improves efficiency . Using the mutually exclusive virtual machine (preferred) policy, the Hadoop cluster can be deployed across different hosts as much as possible, distributing IO pressure and improving overall system performance.
- Example use case for the VMs Affinitive to Hosts (Hard) policy:

Business virtual machines need to be deployed on hosts with a specific CPU frequency to ensure service stability.

- For example: Users deploy four virtual machines to run critical services, requiring hosts with very high CPU frequencies. Since there are few hosts that meet the CPU frequency requirement, the virtual machines must run on these specific hosts. In this case, the virtual machine host affinity (mandatory) policy can be used to make the virtual machines run on the target hosts, ensuring service stability.
- Example use case for the VMs Affinitive to Hosts (Soft) policy:

Different business virtual machines preferably run on hosts located in the same rack to ensure efficient inter-service communication.
- For example: Four virtual machines run different services, which require frequent communication between them, and the hosts need to be as close as possible, such as in the same rack, to minimize communication latency. In this case, the virtual machine host affinity (preferred) policy can be used to make the virtual machines run on the target hosts as much as possible, ensuring efficient inter-service communication.

Advantages

The virtual machine scheduling policies offer the following advantages:

- **Comprehensive & Flexible:**
 - Provides 4 types of scheduling policies and 2 execution mechanisms, defining mutual exclusion/affinity relationships between virtual machines and between virtual machines and hosts. These scheduling policies can be flexibly combined to meet all mainstream business scenario requirements.

- Supports binding/unbinding multiple virtual machines to/from one or more scheduling policies through a virtual machine scheduling group, making the process simple and efficient.
- Displays virtual machine scheduling status intuitively and provides quick conflict resolution operations, allowing users to keep track of business scheduling dynamics in real-time and make adjustments promptly.
- **Powerful & Reliable:**
 - Supports mutual exclusion/affinity between the same/different services, achieving service isolation/efficient intercommunication, ensuring high performance and reliability of services.
 - Allows flexible configuration of virtual machine business fault domains through host scheduling groups, supporting single host deployments, batch host deployments within a single cluster, and cross-cluster host deployments. This avoids single points of failure, ensures business stability, and improves physical resource utilization.

9.5.12.2 Create Mutually Exclusive/Affinitive VM Scheduling Policies

Mutually exclusive/affinitive virtual machine scheduling policies need to be bound to a virtual machine scheduling group to take effect on the virtual machines within the group. Assuming you already have multiple virtual machines running your services, you can follow the steps below to use these scheduling policies:

1. [Create a VM Scheduling Group](#)
2. [Create a Mutually Exclusive/Affinitive VM Scheduling Policy](#)

Create a VM Scheduling Group

In the NexaSphere platform, click **Menu > Business Reliability > VM Scheduling Policy > VM Scheduling Group** to enter the **VM Scheduling Group** page. Click the **New VM Scheduling Group** operation to create a new group.

Set the following parameters:

- **Name:** Set the virtual machine scheduling group name, for example, Virtual Machine Scheduling Group A1
- **Description:** Optional, can be left blank
- **Virtual Machine:** Select one or more target virtual machines to join this scheduling group

After clicking **OK**, the new virtual machine scheduling group will be created successfully.

Create a Mutually Exclusive/Affinitive VM Scheduling Policy

Click **VM Scheduling Policy > New VM Scheduling Policy**, and set the following parameters:

- **Name:** Set the virtual machine scheduling policy name, for example, Virtual Machine Scheduling Policy A1
- **Description:** Optional, can be left blank
- **Type:** Select VM Exclusive from Each Other or VM Affinitive to Each Other
- **Execution Mechanism:** Supports Hard and Soft execution mechanisms
 - **Hard:** Virtual machines are forcibly assigned hosts based on the associated VM scheduling policies. For example, if you associate the VM Exclusive from Each Other policy with a VM scheduling group and select the Hard mechanism for the policy, any two of the virtual machines in the scheduling group are not allowed to run on the same host. If no host is available to be scheduled based on the policy for a virtual machine, the virtual machine will end up failure upon startup.
 - **Soft:** Virtual machines are primarily assigned hosts based on the associated VM scheduling policies. For example, if you associate the VM Exclusive from Each Other policy with a VM scheduling group and select the Soft mechanism for the policy, any two of the virtual machines in the scheduling group will primarily not run on the same host. If no host is available to be scheduled based on the policy for a virtual machine, the virtual machine will attempt to run on a host that does not satisfy the policy.
- **Associate VM Scheduling Group:** After binding, the scheduling policy will take effect on all virtual machines within the virtual machine scheduling group. You can bind existing or newly created scheduling groups. Here, select the scheduling group created above.

After clicking **OK**, the new virtual machine scheduling policy will be created successfully. The virtual machines will run according to the scheduling policy.

9.5.12.3 Create VM Mutually Exclusive/Affinitive Host Scheduling Policies

Virtual machine mutually exclusive/affinitive host scheduling policies need to be bound to both a virtual machine scheduling group and a host scheduling group to take effect on the virtual machines and hosts within the groups. Assuming you have already added multiple hosts running multiple business virtual machines, you can follow the steps below to use these scheduling policies:

1. [Create a VM Scheduling Group](#)

2. [Create a Host Scheduling Group](#)

3. [Create a VM Mutually Exclusive/Affinitive Host Scheduling Policy](#)

Create a VM Scheduling Group

In the NexaSphere platform, click **Menu > Business Reliability > VM Scheduling Policy > VM Scheduling Group** to enter the **VM Scheduling Group** page. Click the **New VM Scheduling Group** operation to create a new group.

Set the following parameters:

- **Name:** Set the virtual machine scheduling group name, for example, Virtual Machine Scheduling Group A2
- **Description:** Optional, can be left blank
- **Virtual Machine:** Select one or more target virtual machines to join this scheduling group

After clicking **OK**, the new virtual machine scheduling group will be created successfully.

Create a Host Scheduling Group

Click **Host Scheduling Group > New Host Scheduling Group**, and set the following parameters:

- **Name:** Set the host scheduling group name, for example, Host Scheduling Group A2
- **Description:** Optional, can be left blank
- **Host:** Select one or more hosts to join this scheduling group

After clicking **OK**, the new host scheduling group will be created successfully.

Create a VM Mutually Exclusive/Affinitive Host Scheduling Policy

Click **VM Scheduling Policy > New VM Scheduling Policy**, and set the following parameters:

- **Name:** Set the virtual machine scheduling policy name, for example, Virtual Machine Scheduling Policy A2
- **Description:** Optional, can be left blank
- **Type:** Select Virtual Machine Mutually Exclusive or affinitive Host
- **Execution Mechanism:** Supports Hard and Soft execution mechanisms
 - **Hard:** Virtual machines are forcibly assigned hosts based on the associated VM scheduling policies. For example, if you associate the VM Exclusive from Each Other policy with a VM scheduling group and select the Hard mechanism for the policy, any two of the virtual machines in the scheduling group are not allowed to run on the same host. If no host is

available to be scheduled based on the policy for a virtual machine, the virtual machine will end up failure upon startup.

- **Soft:** Virtual machines are primarily assigned hosts based on the associated VM scheduling policies. For example, if you associate the VM Exclusive from Each Other policy with a VM scheduling group and select the Soft mechanism for the policy, any two of the virtual machines in the scheduling group will primarily not run on the same host. If no host is available to be scheduled based on the policy for a virtual machine, the virtual machine will attempt to run on a host that does not satisfy the policy.
- **Associate VM Scheduling Group:** After binding, the scheduling policy will take effect on all virtual machines within the virtual machine scheduling group. You can bind existing or newly created scheduling groups. Here, select the scheduling group created above.
- **Associate Host Scheduling Group:** After binding, the scheduling policy will take effect on all hosts within the host scheduling group. You can bind existing or newly created scheduling groups. Here, select the scheduling group created above.

After clicking **OK**, the new virtual machine mutually exclusive/affinitive host scheduling policy will be created successfully. The virtual machines and hosts will run according to the scheduling policy.

9.5.12.4 Manage VM Scheduling Policies and Related Resources

This section uses the Virtual Machine Scheduling Policy A2, Virtual Machine Scheduling Group A2, and Host Scheduling Group A2 as examples to illustrate how to operate virtual machine scheduling policies and related resources.

VM Scheduling Policy

After creating the virtual machine scheduling policy A2, it is enabled by default. You can control the enablement status of the scheduling policy using the **Enable** and **Disable** operations as needed. The execution state of the scheduling policy is associated with its enablement status:

- If the scheduling policy is disabled, then the scheduling state is **Inactive**. In this state, the scheduling policy does not take effect on the associated virtual machines.
- If the scheduling policy is enabled, then the scheduling policy may exist in one of the following two execution states:
 - **Normal:** All virtual machines associated with the virtual machine scheduling policy are running on hosts allocated according to the scheduling policy.

- **Conflict:** Some of the virtual machines associated with the virtual machine scheduling policy are not running on hosts allocated according to the scheduling policy.

If you need to modify basic information about the virtual machine scheduling policy A2, such as the name, description, or execution mechanism, you can click the **Modify Configuration** operation for the scheduling policy, and modify the corresponding information as needed. The changes will be successful upon saving.

If you are sure you no longer need the scheduling policy, you can delete it by clicking the **Actions > Delete** operation for the scheduling policy.

**Note:**

After deleting a virtual machine scheduling policy, the associated virtual machines will no longer be scheduled according to that policy.

VM Scheduling Group

If you want to add more virtual machines to the virtual machine scheduling group A2 for scheduling or remove some virtual machines from the virtual machine scheduling group A2, you can do so in the virtual machine scheduling group main list by clicking the **Add Virtual Machine** or **Remove Virtual Machine** operation for the scheduling group.

If you need to modify basic information about the virtual machine scheduling group A2, such as the name or description, you can click the **Edit Name and Description** operation for the scheduling group, and modify the corresponding information as needed. The changes will be successful upon saving.

If you are sure you no longer need the scheduling group, you can delete it by clicking the **Delete** operation.

**Note:**

After deleting a virtual machine scheduling group, the associated virtual machine scheduling policies will also be deleted. Please proceed with caution.

Host Scheduling Group

If you want to add more hosts to the host scheduling group A2 for scheduling or remove some hosts from the host scheduling group A2, you can do so in the host scheduling group main list by clicking the **Add Host** or **Remove Host** operation for the scheduling group.

If you need to modify basic information about the host scheduling group A2, such as the name or description, you can click the **Edit Name and Description** operation for the scheduling group, and modify the corresponding information as needed. The changes will be successful upon saving.

If you are sure you no longer need the scheduling group, you can delete it by clicking the **Delete** operation.



Note:

After deleting a host scheduling group, the associated virtual machine scheduling policies will also be deleted. Please proceed with caution.

9.5.13 VM HA

The high availability behavior of NexaSphere virtual machines is globally controlled through high availability policies. This section introduces the high availability policy functionality and its usage:

- [Overview](#)
- [HA Policy Basic Operations](#)

9.5.13.1 Overview

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.

High availability policies include the following core concepts:

- Virtual Machine High Availability: Used to set whether virtual machines automatically restart when they are shut down either planned or unexpectedly. If the HA policy is not enabled on the platform, VM HA will take effect after HA policy is enabled.
 - If the high availability switch is turned off: Virtual machines will not automatically restart when they are shut down.
 - If the high availability switch is turned on:
 - Virtual machines will automatically restart when they are shut down plannedly or due to their own unexpected shutdown.
 - If related compute, storage, network, etc., resources experience failures, the virtual machine will migrate to another host and HA start according to a custom-defined fault migration policy as needed.

- **Virtual Machine High Availability Fault Migration Policy:** Used to set whether to migrate a virtual machine to another host to start when related compute, storage, network, etc., resources experience faults.

Fault migration policies support detecting the status of the following resources:

- **Management Network Connection Status:**
 - Detects the network connection status between the host where the virtual machine is located and the management node.
 - If the management node itself fails or the management network is interrupted, it will result in a management network connection status failure.
- **Storage Network Connection Status:**
 - Detects the network connection status between the virtual machine and the data storage resource where its system disk is located.
 - If the data storage where the virtual machine's system disk is located fails or the storage network is interrupted, it will result in a storage network connection status failure for the virtual machine.
- **Business Network Card Status:**
 - If the business network card of the host associated with the distributed switch of the business virtual machine or the network port directly connected to the business network card of the switch fails, it will result in a business network card failure for the virtual machine.

Based on resource status detection, NexaSphere provides four typical fault migration scenarios for easy configuration:

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

Typical Scenario	Management Network Connection Status	Storage Network Connection Status	Business NIC Status	Migrate on Failure?
Scenario D	Failure	Normal	Normal	Do Not Migrate

Use Cases

The following lists typical use cases for high availability policies. If you have similar business scenarios, consider using the high availability policy feature:

- Host Business Network Card Failure Scenario:

You want all associated virtual machines to migrate to another host when the host business network card fails, ensuring high availability for your business.

- For example: Users deploy business virtual machines to host MySQL database services, requiring that virtual machines must not experience extended downtime. You can turn on the high availability switch for these virtual machines and set the business network card status failure to trigger migration. Assuming there are sufficient host resources within the platform, when the business network card of the host where the business virtual machine is located fails, the virtual machine will migrate to another host and start running, without affecting the business operations.
 - Virtual Machine Unexpected Shutdown Scenario:
- You want virtual machines to automatically HA start when they unexpectedly shut down.
- For example: Users deploy business virtual machines to run critical company services, aiming to avoid situations where factors such as host power loss or virtual machine overload cause the virtual machine to shut down and the service cannot automatically recover. You can turn on the high availability switch for these virtual machines. When the virtual machine shuts down, the high availability mechanism will immediately restart it, ensuring business continuity.

Functionality Principles

The high availability policies of NexaSphere mainly include the following two mechanisms:

- Polling to detect the operational status of virtual machines. If a virtual machine shuts down due to its own abnormal condition or planned shutdown, the system checks whether the high

availability switch is turned on. If the switch is on, the virtual machine will be restarted on the current host or another host.

- Polling to detect the status of the host where the virtual machine is located. If any of the management network connection status, storage network connection status, or business network card status is abnormal, the system checks the virtual machine fault migration policy and the virtual machine high availability mode. If the corresponding fault migration switch is turned on and the virtual machine high availability switch is turned on, the virtual machine will migrate to another host and start running.

Benefits of the Functionality

The high availability policies offer the following advantages:

- **Comprehensive & Powerful:** Covers all mainstream high availability scenarios, including various fault scenarios and shutdown scenarios. Ensures the stability and continuity of users' critical business through high availability mechanisms.
- **Flexible & Visual:** Provides an intuitive and simple scenario configuration table, supports one-click configuration of fault migration policies, combines global and virtual machine-level high availability configurations, which can greatly increase the flexibility of business high availability configurations.

9.5.13.2 HA Policy Basic Operations

If you wish to fully understand the basic high availability policy functions of the NexaSphere platform, you can follow these steps:

1. [Enable HA Policy](#)
2. [Set VM Failover Strategy](#)
3. [Set Host Error Detection](#)
4. [Set Advanced Settings](#)
5. [View High Availability Logs](#)
6. [Disable HA Policy](#)

Enable HA Policy

High availability policies in NexaSphere are enabled by default. If they have been disabled, you can click on the **Menu > Business Reliability > HA Policy**, and then turn on the switch at the top of the **HA Policy** page to enable the high availability policies.

Set VM Failover Strategy

After enabling the high availability policies, you can set the high availability migration strategies for the four typical fault scenarios on the **Migration Policies** page:

Typical Scenario	Management Network Connection Status	Storage Network Connection Status	Business NIC Status	Migrate on Failure?	Migration Explanation
Scenario A	Normal	Normal	Failure	Migrate Do Not Migrate	Supports setting to migrate or do not migrate.
Scenario B	Normal	Failure	Normal	Migrate Do Not Migrate	Supports setting to migrate or do not migrate. However, in a SAN storage environment, if set to do not migrate here, the storage network connection status failure will still trigger automatic migration.
Scenario C	Normal	Failure	Failure	Migrate Do Not Migrate	The migration strategy for when both the storage connection status and the business network card status fail follows the migration strategy for either status failing: - If the migration strategy for both the storage connection status and business network card status failure scenarios is set to do not migrate, then this is set to do not migrate. - If the migration strategy for one of the failure scenarios is set to migrate, then this is set to migrate.
Scenario D	Failure	Normal	Normal	Do Not Migrate	When the management network status is faulty, it is not supported to set a fault migration strategy.

**Note:**

Storage network connection status only supports detecting shared storage and does not currently support local storage.

Set Host Error Detection

After enabling the high availability policies, you can set the host failure judgment policies on the **Migration Policies** page:

Host Error Detection Item	Description
Host Self-Inspection Interval	The interval that a host inspects its own status. Default: 5. Unit: second.
Maximum Host Self-Inspection Attempts	The maximum number of attempts that a host inspects its own status . If the self-inspection of a host fails by the maximum attempts, it is determined that network errors occur with the host. Default: 6.

Set Advanced Settings

After enabling the high availability policies, you can set the advanced settings for high availability policies on the **Advanced Settings** page, including advanced settings for both virtual machines and hosts.

Category	Name	Description
Virtual Machine	HA VM State Update Speed	The speed of updating the state of NeverStop virtual machines on the UI. Default: 1. Valid values: -1 to 5. A higher value indicates a lower update speed. However, a lower update speed makes the system ignore a lot of outdated notifications, thus decreasing the system workload. If set to -1, the NeverStop VM states on the UI are not updated automatically.
	Maximum Interval for VM Attempt to HA Start	The maximum interval for the system to finish the GC (garbage collection) job and attempt to restart a NeverStop VM according to the HA policy after the virtual machine is stopped unexpectedly. Default: 300. Unit: second.
	VM Retry HA Start Interval	The interval for a Neverstop VM to retry an HA start after the previous HA start attempt fails. Default: 60. Unit: second.

Category	Name	Description
	HA VM State Scanning Interval	The interval to scan the status of a NeverStop VM after it fails to HA start. Default: 60. Unit: second.
Host	Timeout Period for Host Connecting to Data Storage	The time for hosts to attempt to connect to data storage. If a host fails to connect to a data storage during this period, its connection attempt is determined as timeout. Default: 5. Unit: second.
	Abnormal Host Status Update Interval	The interval for the system to check and update the status of abnormal hosts. Default: 5. Unit: second.
	Minimum Connection Attempts Required to Determine Host is Disconnected	The maximum times for the system to attempt to connect to a host. If the system fails to connect to the host after the specified times of attempt, the host is determined as disconnected. Default: 12.
	Ping Response Time to Determine Host Connection is Established Successfully	The time period for the system to wait the host response after it pings the host. Receiving a response within this period indicates that the system establishes a successful connection with the host. Default: 5. Unit: second.
	Minimum Connection Success Rate to Determine Host is Re-Connected	The minimum rate of successful connections occupied in total connection attempts to determine a disconnected host is successfully re-connected. Default: 50. Unit: %.
	Minimum Successful Connections Required to Determine Host is Re-Connected	The minimum successful connections that the system has to establish with a disconnected host before the host can be determined as re-connected. Default: 5.

View High Availability Logs

After enabling the high availability policy, if the platform triggers the high availability mechanism, high availability logs will be generated. You can view these logs on the **O&M Management > Tasks > HA Task** page. The logs support viewing task results, virtual machine names, virtual machine owners, previous hosts, target hosts, start times, and completion times, enriching operational scenarios for auditing and tracing.

- Supports selecting a time period to view high availability logs for virtual machines during the selected period. Available time periods include: last 7 days, last month. By default, the latest 7 days of logs are displayed.
- Supports custom time periods to view high availability logs for virtual machines during the set period.
- Supports searching for high availability logs for virtual machines by entering the virtual machine name or owner.
- Supports filtering high availability logs for virtual machines by task result. Task results include: success, failure.
- Supports sorting high availability logs for virtual machines by start/completion time.
- Supports exporting high availability logs for virtual machines in CSV format.
- Supports adjusting the number of completed high availability logs for virtual machines displayed per page. Selectable values are: 10, 20, 50, 100, and pagination is supported.

Disable HA Policy

If you wish to globally disable the high availability feature for virtual machines, you can do so on the **HA Policy** page by clicking the **Disable** action.



Note:

After disabling the high availability policy, virtual machines will not automatically restart upon shutdown, which may cause service interruptions. Proceed with caution.

9.5.13.3 Implement HA Policy in Business Practices

Assume you have deployed four business virtual machines on Host A to support MySQL database services. To ensure high availability, if Host A's service network adapter fails, all four virtual machines should be migrated to another host. In this scenario, set the high availability mode for these virtual machines to **NeverStop**, configure the policy to trigger a **migration** when a service network adapter failure occurs, and ensure there are sufficient resources available on other hosts within the platform.

To configure the high availability policy for this scenario, follow these steps:

1. Enable the high availability policy: In the NexaSphere platform, click **Menu > Business Reliability > HA Policy**. Turn on the switch at the top of the **HA Policy** page.
2. Turn on the high availability switch for virtual machines:

You can set this up in two ways, with precedence order being: virtual machine level > cluster level.

- When creating a new virtual machine, turn on the **HA** switch to enable it.
 - Go to the cluster page where the virtual machine resides, then click **Modify Configuration > Advanced Settings > VM Settings**. Turn on the **VM HA** switch. This ensures that any new virtual machines created in this cluster will have the high availability switch turned on by default.
3. Configure the VM fault migration policy: Navigate to the **HA Policy > Migrate Strategy** page and turn on the switch corresponding to **Fail Over** under **Scenario A**. When this switch is enabled, the **Fail Over** switch under **Scenario C** will automatically turn on as well.

After configuring the high availability policy for virtual machines, if Host A's service network adapter fails, the four virtual machines on that host will immediately migrate and start on Host B. You can search for related migration logs in **O&M Management > Tasks > HA Task**.

9.5.14 VM Security

9.5.14.1 Configure and Manage a Standard Key Provider

9.5.14.1.1 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, set the following parameters:
 - **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.
- **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 [Establish Trust](#).
- Scenario practices: [Use the Thales Key Provider](#), [Use the HashiCorp Key Provider](#)

9.5.14.1.2 Establish Trust

9.5.14.1.2.1 Make KMS Trust NexaSphere

9.5.14.1.2.1.1 Use KMS Certificate and Private Key to Make KMS Trust

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](#).

9.5.14.1.2.1.2 Use the New Certificate Signing Request (CSR) to Make KMS Trust

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](#).

9.5.14.1.2.2 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**.

9.5.14.1.3 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, perform the **Update Key** operation on the relevant virtual machines before deletion, and re-encrypt them using the default key provider.

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, all resources encrypted with this key provider (such as TPM-enabled virtual machines, encrypted virtual machines, and encrypted disks) will enter a locked state and become inaccessible.

5. Read the risk warning, enter the confirmation information, and click **OK**.

9.5.14.1.4 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 9-6: Add a User




- c) Create the localhost user.

Figure 9-7: 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 9-8: 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 KMS Trust the Platform**.
 - d) In the **Make KMS Trust the Platform** 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 9-9: Access the KMIP Module


- i) Click **Client Profile > Add Profile**.

Figure 9-10: Add a Profile


- j) In the **Add Profile** dialog, set the profile basic information, paste the CSR information, and then click **Save**.

Figure 9-13: Create a Token


m) Click **Registered Clients > Add Client**.

Figure 9-14: Add a Client

n) In the **Add Client** dialog, set the client name, paste the registration token, and then click **Save**.

Figure 9-15: Configure the Client


- o) After saving successfully, click **Save CSR** and **Save Certificate** to download the CSR and certificate.

Figure 9-16: 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 9-17: Access the KMIP Module




- c) Click **Client Profile > Add Profile**.

Figure 9-18: Add a Profile




d) In the **Add Profile** dialog, set the profile name, and then click **Save**.

Figure 9-19: Add a Profile




e) Click **Registration Token > New Registration Token**.

Figure 9-20: 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 9-21: Create a Token




- g) Click **Registered Clients > Add Client**.

Figure 9-22: Add a Client




- h) In the **Add Client** dialog, set the client name, paste the registration token, and then click **Save**.

Figure 9-23: Configure the Client



- i) After saving successfully, click **Save Certificate** and **Save Private Key** to download the certificate and private key.

Figure 9-24: 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 KMS Trust the Platform**.
- m) In the **Make KMS Trust the Platform** 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**.

9.5.14.1.5 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 9-25: Access the KMIP Module



d) On the **kmip** page, click **Create scope** to create a scope or enter an existing one.

Figure 9-26: 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 9-27: 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 9-28: Generate a Certificate



- i) Copy the certificate and private key.

Figure 9-29: 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 KMS Trust the Platform**.
 - m) In the **Make KMS Trust the Platform** 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.

9.5.14.2 Configure and Manage a Native Key Provider

9.5.14.2.1 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, set the following parameters:
 - **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](#).

9.5.14.2.2 Back Up the Native Key Provider

The native key provider is stored only in NexaSphere. If NexaSphere fails or needs to be rebuilt, the key is lost and all encrypted virtual machines cannot be decrypted or started. Therefore, backing up the native key provider is necessary for disaster recovery. After you back it up, you can restore key services and remount encrypted virtual machines after you rebuild 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.

- **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 configuration in case of disaster.

If you forget or lose the password:

- You cannot restore the key provider.
- You cannot access encrypted resources (such as TPM-enabled virtual machines and encrypted virtual machines) that depend on this key provider.

7. Click **OK**.

The backup file is downloaded through your browser. Save the backup file in a secure location.

9.5.14.2.3 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**.

9.5.14.2.4 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, perform the **Update Key** operation on the relevant virtual machines before deletion, and re-encrypt them using the default key provider.

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, all resources encrypted with this key provider (such as TPM-enabled virtual machines, encrypted virtual machines, and encrypted disks) will enter a locked state and become inaccessible.
- 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**.

9.5.14.3 Rekey

For business or compliance requirements (such as regular key rotation, suspected key leakage, or meeting security audit standards), you can update the key and re-encrypt encrypted virtual machines using the default key provider.

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**.

This operation re-encrypts all encrypted resources using the new key from the current default key provider.

9.5.14.4 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 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 New 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**.

10 Storage Management

This chapter mainly introduces how to use storage virtualization resources, including data storage.

This section covers the use of data storage in the following topics:

- [Add a Data Storage](#)
- [Modify Data Storage Configuration](#)
- [Data Storage Cleanup & Deletion](#)

10.1 Add a Data Storage

10.1.1 Add a Local Storage

If you wish to use the local hard disk directory of a host to create storage resources, you can refer to this section for instructions.

Prerequisites

- Make sure to plan the mount path for the local storage on the host's disk in advance.
- Check the quantity and type limitations between the cluster and the data storage to the attached. For more information, see [Cluster 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, set the following parameters:

Basic Information

- **Name:** Name of the local storage.
- **Description:** Optional. You can add relevant information as a note.
- **Type:** Select **Local Storage**.
- **Data Center:** Location of the data center where the local storage resides.

Configurations

- **Cluster:** The cluster where the local storage needs to be attached.
- **Addition Method:** Supports **Free Disk** and **Local Directory** methods

If you choose **Free Disk**, you need to configure the following parameter:

- **Host Disk:** Add unmounted or unpartitioned free disks on the host

**Note:**

Configuring a free disk will format the selected disk, completely clearing all partitions, file systems, and data on the disk.

- **Mount Path:** The mount path for the local storage on the host's disks

**Note:**

System directories such as `/`, `/dev/`, `/proc/`, `/sys/`, `/usr/bin`, and `/bin` cannot be used. Using system directories might cause the hosts unable to work properly.

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

10.1.2 Add an NFS Storage

If you wish to use a Network File System (NFS) to create storage resources, you can refer to this section for instructions.

Prerequisites

- Make sure to plan the mount path and set the appropriate directory permissions on the NFS Server in advance.
- Check whether the NFS Server supports mount parameters in advance. If it does, you can specify relevant parameters when adding NFS storage to optimize the performance, security, and reliability of the network file system.
- It is recommended to plan a separate storage network in advance to avoid network congestion.
- Check the quantity and type limitations between the cluster and the data storage to the attached. For more information, see [Cluster 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, set the following parameters:

Basic Information

- **Name:** Name of the NFS storage.
- **Description:** Optional. You can add relevant information as a note.
- **Type:** Select **NFS**.
- **Data Center:** Location of the data center where the NFS storage resides.

Configurations

- **Cluster:** The cluster where the NFS storage needs to be attached.
- **Mount Path:** The shared directory URL of the NFS Server. 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 the hosts unable to work properly.

- **Mount Option:** Mount parameters for the NFS Server end, you can refer to the content in the *-o* parameter of *mount*. If the parameters set here conflict with those on the NFS Server, the NFS Server's settings take precedence.
- **Storage Network:** The CIDR of the storage network specified for the NFS storage.

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

10.1.3 Add a SAN Storage

If you wish to use shared LUNs to create storage resources, you can refer to this section for instructions.

Prerequisites

- LUNs are provided by iSCSI, FC, or NVMe storage. For more information, see .
- Check the quantity and type limitations between the cluster and the data storage to the attached. For more information, see [Cluster 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, set the following parameters:

Basic Information

- **Name:** Name of the SAN storage.
- **Description:** Optional. You can add relevant information as a note.
- **Type:** Select **SAN Storage**.
- **Data Center:** Location of the data center where the SAN storage resides.

Configurations

- **Cluster:** The cluster where the SAN storage needs to be attached.
- **Default Provisioning Type:** The allocation method for disk storage space. Default: Thick Provision. Options include Thick Provision and Thin Provision.
 - Thick Provision: Pre-allocate the required storage space, providing sufficient storage capacity to ensure storage performance.
 - Thin Provision: Allocate storage space according to actual usage, achieving higher storage utilization.
- **Storage Network:** Specify a storage network for the shared storage. The system uses the storage network to check the health status of virtual machines.
- **Storage Addition Method:** Select a method to add the storage.
 - New: Creates a brand new data storage by initializing an unused LUN into a data storage usable by the platform. You can choose whether to clear existing data on the LUN device during addition.
 - Register: Imports LUNs that already contain data (for example, storage migrated from another platform or replicated through disaster recovery) into the platform. The platform will identify its SAN storage structure and repair the metadata. During registration, all LUN devices under that SAN storage will be automatically added.
- **LUN:** Select LUN devices provided by iSCSI storage, FC storage, or NVMe storage.
- **Cleanse LUN:** Choose whether to forcibly clear data from the LUN devices, such as file system, RAID, or partition table labels. Default: Unselected.

**Note:**

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

- **Storage UUID:** When you select **Register** as the storage addition method, you must specify how to handle the storage UUID.
 - Keep: Retains the existing UUID of the SAN storage unchanged.
 - Reset: Regenerates a new UUID for the SAN storage. This is useful in scenarios such as UUID conflicts after cloning or snapshots, or when decoupling from the production environment during disaster recovery drills. After reset, the original UUID becomes invalid. If other hosts continue to access the storage using the original UUID, I/O errors will occur.

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

10.1.3.1 Add an iSCSI Storage

iSCSI storage is a SAN storage that uses the iSCSI protocol for data transmission.

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, set the following parameters:
 - **Name**: Name of the iSCSI storage.
 - **IP Address**: IP address of the iSCSI storage server.
 - **Port**: Target port for the iSCSI storage. Default: 3260.
 - **Cluster**: The cluster where the iSCSI storage will be attached. You can also attach the iSCSI storage to a cluster after the addition.



Note:

If you want to use the LUNs provided by the iSCSI storage, you need to attach the iSCSI storage to the cluster where the SAN storage is attached.

- **CHAP Username**: CHAP authentication username.
 - **CHAP Password**: CHAP authentication password.
6. Review the configuration and click **OK**.

What's next

You can add iSCSI LUNs as SAN storage or pass through them to a virtual machine. LUNs that have been added as SAN storage cannot be passed through to a virtual machine.

- To add iSCSI LUNs as SAN storage, see [Add a SAN Storage](#).
- To pass through LUNs to a virtual machine through RDM disks, see [Create a New Virtual Machine](#).

10.1.3.2 Synchronize a FC Storage

FC storage is a SAN storage that uses the FC protocol for data transmission.

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 through them to a virtual machine. LUNs that have been added as SAN storage cannot be passed through to a virtual machine.

- To add FC LUNs as SAN storage, see [Add a SAN Storage](#).
- To pass through LUNs to a virtual machine through RDM disks, see [Create a New Virtual Machine](#).

10.1.3.3 Add a NVMe Storage

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, set the following parameters:
 - **Name**: Name of the NVMe storage.
 - **Transmission Mode**: Support two transmission modes: RDMA and TCP.
 - **IP Address**: IP address of the NVMe storage.
 - **Port**: Port of the NVMe storage.
 - **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 through them to a virtual machine. LUNs that have been added as SAN storage cannot be passed through to a virtual machine.

- To add NVMe LUNs as SAN storage, see [Add a SAN Storage](#).
- To pass through LUNs to a virtual machine through RDM disks, see [Create a New Virtual Machine](#).

10.1.4 Add a nSDS Distributed Storage

If you wish to use distributed block storage to create storage resources, you can refer to this section for instructions.

Prerequisites

- You have set up the nSDS distributed storage on the storage side. If you wish to use the specified image cache pool and storage pool, you need to create the corresponding storage pools in the distributed storage cluster in advance.
- It is recommended to plan a separate storage network in advance to avoid network congestion.
- Check the quantity and type limitations between the cluster and the data storage to the attached. For more information, see [Cluster 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, set the following parameters:

Basic Information

- **Name:** Name of the distributed storage.
- **Description:** Optional. You can add relevant information as a note.
- **Type:** Select **Distributed Storage nSDS**.
- **Data Center:** Location of the data center where the nSDS distributed storage resides.

Configurations

- **Cluster:** The cluster where the nSDS distributed storage needs to be attached.
- **Key Authentication:** Choose whether to use key authentication for the nSDS distributed storage. Default: Enabled.

**Note:**

Make sure that the authentication option on the distributed storage side is consistent with this option. If the authentication is disabled on the nSDS distributed storage side but enabled here, you might fail to create virtual machines, and vice versa.

- **Monitoring Node:** Add monitoring nodes and complete the monitoring node IP, SSH port, username, and password configurations.
- **Image Cache Pool:** Specify a storage pool for image caches. If you do not specify a storage pool, the system creates one automatically.
- **Storage Pool:** Specify a storage pool for data disks. If you do not specify a storage pool, the system creates one automatically.
- **Storage Network:** The CIDR of the storage network specified for the nSDS distributed storage.

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

10.1.5 Add a ZHPS Distributed Storage

If you wish to use vhost-user mode to connect with high-performance SSD distributed storage, you can refer to this section for instructions.

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 the attached. For more information, see [Cluster and Data Storage](#).
- Make sure the platform is installed with 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, set the following parameters:

Basic Information

- **Name:** Name of the high-performance storage.
- **Description:** Optional. You can add relevant information as a note.
- **Type:** Select **Distributed Storage ZHPS**.
- **Data Center:** Location of the data center where the ZHPS distributed storage resides.

Configurations

- **Cluster:** The cluster where the ZHPS distributed storage needs to be attached.
- **IP Address:** The management address of the ZHPS distributed storage.
- **Port:** The port corresponding to the management address of the ZHPS distributed storage.
- **Username:** The username of the management address of the ZHPS distributed storage.
- **Password:** Password corresponding to the username.
- **Storage Pool:** After a successful connection test, specify storage pools. You need to create storage pools on the storage side in advance.

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

10.1.6 Add a ZBS Distributed Storage

If you wish to Connects high-performance distributed block storage through the CBD interface, you can refer to this section for instructions.

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 the attached. For more information, see [Cluster and Data Storage](#).
- Make sure the platform is installed with 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, set the following parameters:

Basic Information

- **Name:** Name of the ZBS distributed storage.
- **Description:** Optional. You can add relevant information as a note.
- **Type:** Select **Distributed Storage ZBS**.
- **Data Center:** Location of the data center where the ZBS distributed storage resides.

Configurations

- **Cluster:** The cluster where the ZBS distributed storage needs to be attached.
- **MDS Node:** Add MDS nodes. You need to specify the MSD node management IP, SSH port, username, and password.

- **Storage Pool:** Enter the pool name after creating the storage pool on the storage side in advance.

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

10.2 Modify Data Storage Configuration

If you have added data storage, you can modify the data storage configuration as needed based on your business scenarios.

- [Modify Local Storage Configuration](#)
- [Modify NFS Storage Configuration](#)
- [Modify SAN Storage Configuration](#)
- [Modify Distributed Storage Configuration](#)

10.2.1 Modify Local Storage Configuration

If you have added local storage, you can modify the local storage configuration as needed based on your business scenarios, including basic and advanced configurations.

- [Modify Basic Configuration](#)
- [Modify Advanced Configuration](#)

Modify Basic Configuration

If you only need to modify the name and description of the local storage, you can click **Action > Edit Name and Description** on the target local storage page to make the changes.

Modify Advanced Configuration

If you want to modify the advanced configuration information for local storage, including the storage over-provisioning ratio and hard disk pre-allocation strategy, you can follow these steps:

1. Navigate to the target local storage page.
2. Enter the **Advanced Settings** sub-page.
3. Click **Edit**.

NexaSphere supports modifying the following advanced configurations for local storage:

- **Data Storage Overcommit Ratio:** Used to control the allocatable space of data storage for virtual machines.

Calculation formula: Data storage allocatable capacity = [(actual capacity - reserved capacity) × over-provisioning ratio] - (threshold capacity + sum of all hard disks allocated to virtual machines + snapshots + image cache + migration cache)

- **Disk Preallocation Policy:** Used to set the hard disk pre-allocation strategy in data storage. The default is none.
 - none: Does not use pre-allocation strategy. With this strategy, files are dynamically allocated the required space as needed.
 - metadata: Only pre-allocates the space required for file metadata, without allocating any space for the data. With this strategy, the required storage space for the hard disk is dynamically allocated as data is written.
 - falloc: Pre-allocates the required space for the file but does not immediately erase the corresponding data on the physical device. Erasure occurs when the virtual machine first writes data to the hard disk.
 - full: Pre-allocates the required space for the file.

10.2.2 Modify NFS Storage Configuration

If you have added NFS storage, you can modify the NFS storage configuration as needed based on your business scenarios, including basic and advanced configurations.

- [Modify Basic Configuration](#)
- [Modify Advanced Configuration](#)

Modify Basic Configuration

If you only need to modify the name and description of the NFS storage, you can click **Action > Edit Name and Description** on the target NFS storage page to make the changes.

Modify Advanced Configuration

If you want to modify the advanced configuration information for NFS storage, including the storage over-provisioning ratio and hard disk pre-allocation strategy, you can follow these steps:

1. Navigate to the target NFS storage page.
2. Enter the **Advanced Settings** sub-page.
3. Click **Edit**.

NexaSphere supports modifying the following advanced configurations for NFS storage:

- **Data Storage Overcommit Ratio:** Used to control the allocatable space of data storage for virtual machines.

Calculation formula: Data storage allocatable capacity = [(actual capacity - reserved capacity) × over-provisioning ratio] - (threshold capacity + sum of all hard disks allocated to virtual machines + snapshots + image cache + migration cache)

- **NFS Storage Disk Preallocation Policy:** Used to set the hard disk pre-allocation strategy in NFS storage. The default is none.
 - none: Does not use pre-allocation strategy. With this strategy, files are dynamically allocated the required space as needed.
 - metadata: Only pre-allocates the space required for file metadata, without allocating any space for the data. With this strategy, the required storage space for the hard disk is dynamically allocated as data is written.
 - falloc: Pre-allocates the required space for the file but does not immediately erase the corresponding data on the physical device. Erasure occurs when the virtual machine first writes data to the hard disk.
 - full: Pre-allocates the required space for the file.

10.2.3 Modify SAN Storage Configuration

If you have added SAN storage, you can modify the SAN storage configuration as needed based on your business scenarios, including basic configuration, advanced configuration, and expanding the SAN storage.

- [Modify Basic Configuration](#)
- [Modify Advanced Configuration](#)
- [Expand SAN Storage](#)

Modify Basic Configuration

If you only need to modify the name and description of the SAN storage, you can click **Action > Edit Name and Description** on the target SAN storage page to make the changes.

Modify Advanced Configuration

If you want to modify the advanced configuration information for SAN storage, including the hard disk pre-allocation strategy, storage allocation strategy, and storage over-provisioning ratio, you can follow these steps:

1. Navigate to the target SAN storage page.

2. Enter the **Advanced Settings** sub-page.
3. Click **Edit**.

NexaSphere supports modifying the following advanced configurations for SAN storage:

- **Disk Preallocation Policy:** Used to set the hard disk pre-allocation strategy in SAN storage. The default is metadata.
 - metadata: Only pre-allocates the space required for hard disk metadata, without allocating any space for the data. With this strategy, the required storage space for the hard disk is dynamically allocated as data is written.
 - none: Does not use pre-allocation strategy.
- **SAN Storage Allocation Strategy:** Used to set the landing strategy for hard disks and snapshots on SAN storage LUN devices. The default is according to system allocation. Available strategies include according to system allocation, hard disks created in the LUN with the most remaining capacity, and hard disks created in the LUN with the least number of LVs (snapshots + hard disks).
- **Data Storage Overcommit Ratio:** Used to control the allocatable space of data storage for virtual machines.

Calculation formula: Data storage allocatable capacity = [(actual capacity - reserved capacity) × over-provisioning ratio] - (threshold capacity + sum of all hard disks allocated to virtual machines + snapshots + image cache + migration cache)

Expand SAN Storage

If during the use of SAN storage, you find that the storage capacity is insufficient to meet your business needs, you can add new LUN devices or expand existing LUN devices. You can follow these steps to add LUN devices:

1. In the physical environment, expand the LUN device for iSCSI storage, FC storage, or NVMe storage.

After successfully expanding the LUN device, return to the platform. You can view the usage of block devices on the **Data Center > Storage Target** sub-page. You can also manually click the refresh button to get the latest storage information.

2. Navigate to the target SAN storage page.
3. Enter the **Shared Block** sub-page.
4. Click **Add** and select the newly expanded LUN device.

**Note:**

If the LUN device contains data, it may cause failure when adding the LUN device or mounting the SAN storage. You can check **Cleanup Block Device** on the **Add Shared Block** page.

If you need to obtain the latest capacity information for a LUN device, you can do so on the **Shared Block** sub-page by selecting the target LUN device and clicking **Refresh Capacity**.

10.2.4 Modify Distributed Storage Configuration

If you have added distributed storage, you can modify the distributed storage configuration as needed based on your business scenarios, including basic configuration, advanced configuration, adding, deleting, or modifying monitoring nodes, and adding, deleting, or modifying storage pools.

- [Modify Basic Configuration](#)
- [Modify Advanced Configuration](#)
- [Modify Monitoring Nodes](#)
- [Modify Storage Pool](#)

Modify Basic Configuration

If you only need to modify the name and description of the distributed storage, you can click **Action > Edit Name and Description** on the target distributed storage page to make the changes.

Modify Advanced Configuration

If you want to modify the advanced configuration information for distributed storage, including the storage over-provisioning ratio, you can follow these steps:

1. Navigate to the target distributed storage page.
2. Enter the **Advanced Settings** sub-page.
3. Click **Edit**.

NexaSphere supports modifying the following advanced configurations for distributed storage:

- **Data Storage Overcommit Ratio:** Used to control the allocatable space of data storage for virtual machines.

Calculation formula: Data storage allocatable capacity = [(actual capacity - reserved capacity) × over-provisioning ratio] - (threshold capacity + sum of all hard disks allocated to virtual machines + snapshots + image cache + migration cache)

Modify Monitoring Nodes

Add Monitoring Node

You can follow these steps to add a monitoring node to the distributed storage:

1. Navigate to the target distributed storage page.
2. Click **Monitoring Nodes > Add Monitoring Node**.
 - **Mon Node Management IP**: The IP address of the monitoring node
 - **SSH Port**: The SSH port of the monitoring node
 - **Username**: The SSH username for the monitoring node
 - **Password**: The password for the SSH username of the monitoring node
3. After confirming that the configuration information is correct, click **OK** to add a monitoring node.

Modify Monitoring Node Configuration

If you want to adjust the configuration of an added monitoring node, including modifying the SSH username, SSH password, SSH port, and Mon port, you can do so on the **Monitoring Nodes** sub-page by clicking **Actions** and making the necessary modifications.

Delete Monitoring Node

If you need to delete an existing monitoring node, you can do so on the **Monitoring Nodes** sub-page by clicking **Actions > Delete Monitoring Node**. You can also delete monitoring nodes in batch.



Note:

Deleting a monitoring node may cause the distributed storage cluster to lose connectivity. Please proceed with caution.

Modify Storage Pool

Add Storage Pool

You can follow these steps to add a storage pool to the distributed storage:

1. Navigate to the target distributed storage page.
2. Click **Storage Pools > Add Storage Pool**.
 - **Pool Name**: The UUID of the storage pool
 - **Display Name**: A custom display name for the storage pool
3. After confirming that the configuration information is correct, click **OK** to add a storage pool.

Modify Storage Pool Configuration

If you want to edit the display name of a storage pool, you can do so on the **Storage Pools** sub-page by clicking **Actions > Set Display Name** and making the necessary changes.

Delete Storage Pool

If you need to delete an existing storage pool, you can do so on the **Storage Pools** sub-page by clicking **Actions > Delete**.

**Note:**

- The image cache pool does not support deletion operations.
- If there is data in the storage pool, it cannot be deleted.
- You must retain at least one storage pool.

10.3 Data Storage Cleanup & Deletion

You can refer to the content in this section to clean up data storage space or delete data storage as needed.

- [Clean up Data Storage](#)
- [Delete Data Storage](#)

Clean up Data Storage

You can also clean up the original data retained in the data storage due to cross-storage migration. On the **Data Cleanup** sub-page of the target data storage details page, click **Action > Cleanup** to clean up the original data.

**Note:**

Make sure that the migrated data is intact after the storage migration. The original data cannot be recovered after cleanup. Please proceed with caution.

Delete Data Storage

If you no longer need a particular data storage, you can delete it to save storage resources.

Before deleting data storage, you need to unload it from the cluster, otherwise, deletion will not be possible.

You can delete data storage through the following methods:

- Delete a single data storage: Navigate to the target data storage page and click **Actions** > **Delete** to perform the deletion operation.
- Bulk delete data storage: Navigate to the **Data Center** > **Data Storage** sub-page, select the data storages you want to delete. Then click **Bulk Action** > **Delete** to perform the deletion operation.

**Note:**

During deletion, all resources on the selected data storage, including virtual machines, hard disks, snapshots, etc., will be deleted. Please proceed with caution.

11 Network Management

This chapter mainly introduces how to use network virtualization resources and services, including distributed switches, distributed port groups, and security groups. This section covers the following topics on how to use network resources and services:

- [Network Resource](#)
- [Network Service](#)

11.1 Network Resource

11.1.1 Distributed Switch

Through distributed switches, you can set up and configure network connections in the NexaSphere environment.

After the first host is added to the cluster, NexaSphere automatically creates a default distributed switch, a default distributed port group, and a default Kernel adapter based on the host's related configuration. These are used for centralized management of the host's management network. Based on your network planning, you can flexibly reuse the default distributed switch or use a newly created distributed switch.

11.1.1.1 Create a Distributed Switch

Procedure

1. In the navigation pane, choose **Inventory > Network Resource**.
2. In the resource tree, right-click a target data center and select **New Distributed Switch**.
3. In the **New Distributed Switch** dialog:
 - a) Complete the basic information configuration.
 - **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.
 - **Cluster**: Select the target cluster for the distributed switch (virtualization cluster or bare metal cluster).
 - b) Complete the network configuration.

This step only appears when you select a virtualization cluster.

- **Addition Method:** Supports individual addition, batch aggregation, and specifying same network interface.

When selecting individual addition or batch aggregation, configure these parameters:

- **Uplink Name:** Set a name for aggregated host physical ports connecting to physical switches.



Note:

- By default, naming follows the format "Uplink + suffix," with the suffix automatically incrementing as "1/2/3/..." to distinguish between resources. When the number of uplinks is 10 or more, the default naming format changes to "Up + suffix".
 - Custom uplink name must be within 1 to 10 characters and can only include English letters, numbers, and the special characters "-" and "_". The name cannot start with a number.
- **Bond Mode:** Select a bonding mode for physical ports.
 - LACP (mode 4): Bonded ports share the same speed and duplex settings. Network traffic is evenly distributed across all ports for load balancing. This mode supports 1 to 8 physical ports. We recommend bonding at least 2 ports.
 - Active-Backup (mode 1): Bonded ports work in active-backup mode. Normally, the active port handles all traffic. If active port fails, the backup port takes over automatically. This mode supports 1 to 8 physical ports. We recommend bonding 2 ports.
 - **Hash Policy:** When selecting LACP mode, you can configure the hash policy to determine network traffic egress.
 - 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:** Select host ports to be bonded.
 - When creating host bonded ports individually, all selected ports on the same host must have the same speed.

- When creating host bonded ports in bulk, you can only select ports with identical speeds.

When selecting specifying same network interface, configure these parameters:

- **Network Interface Type:** Select the interface type, including Aggregated Interface and Non-Aggregated Interface.
- **Bond:** Select ports to all hosts in the cluster that have matching ports.

c) Complete the distributed port group configuration.

By default, the New Distributed Port Group checkbox is selected. You can choose whether to create a distributed port group on this distributed switch.

- **Name:** Enter a name for the distributed port group.
- **VLAN Type:** Select a VLAN type. When selecting Standard VLAN, you need to specify a VLAN ID.
- **DHCP Service:** Choose whether to enable the automatic IP address assignment for platform resources.
- **IP Address Management:** Choose whether to enable the IP address management. When enabled, you can add network ranges to this distributed port group. IP addresses in these ranges can be allocated via DHCP service (when enabled) to resources in the network.
- **IP Address Type:** Supports IPv4 and IPv6.

When selecting IPv4, configure the following parameters:

- **IP Allocation Policy:** After enabling the DHCP service, you can select one of the following policies to assign IP addresses.
 - Random Allocation: The system randomly assigns IP addresses from the network range.
 - Allocate in Order: The system assigns all available IP addresses from the 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 available IP addresses from the network range in ascending order. Released IP addresses are assigned when currently available IP addresses are used up.

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:** Supports IP range and CIDR.
 - For IP range, enter start IP, end IP, netmask, and gateway.

**Note:**

Do not include gateway, broadcast address, or network addresses in the IP range.

- For CIDR, enter CIDR Block and gateway. You can enter the first or last CIDR address for gateway. If left blank, the first CIDR address will be used.

When selecting IPv6, configure the following parameters:

- **Network Range Method:** Supports IP range and CIDR.
 - For IP range, enter start IP, end IP, prefix length, and gateway.
 - For CIDR, enter CIDR Block.
- **IP Configuration Mode:** Select IPv6 address allocation method.
 - **Stateful-DHCP (Default):** The interface address and other parameters are all configured through DHCP. The IP range method supports stateful DHCP.
 - **Stateless-DHCP:** The interface address is automatically derived from the route advertisement prefix and the interface Mac address. Other parameters are configured through DHCP.
 - **SLAAC:** The interface address is automatically derived from the prefix of the route advertisement that also contains other parameters.
- **DHCP IP:** Specify an IP address used by the DHCP service.
- **DNS:** Specify a DNS to provide DNS resolution service for the distributed port group.

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

**Note:**

- When adding either a single NIC or a bonded NIC to the distributed switch, the bridge names created by distributed port groups will all start with `br_dvs`.
- If the management network address was originally configured on a physical NIC or sub-interface before adding the host, it will be moved to the `br_dvs{ID}_{VLAN ID}` bridge after addition. When removing the host, the management address will revert back to the physical NIC or sub-interface.
- If the management network address was already on a `br` bridge (for example, a user-defined `br_{name}_{VLAN ID}`) before adding the host, the bridge name will remain unchanged. In scenarios where management and business networks share NICs, when the new distributed port group's VLAN differs from the management network VLAN, the created bridge will still follow the `br_dvs{ID}_{VLAN ID}` naming convention, but this will not affect the NIC hosting the management address.

11.1.1.2 Distributed Switch Uplink

11.1.1.2.1 Modify Uplink Configurations

Procedure

1. In the navigation pane, choose **Inventory > Network Resource**.
2. In the resource tree, select a target distributed switch.
3. On the target distributed switch's 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 there are joined hosts, but only if the default distributed switch uses a single NIC.
- When no hosts are joined, you can only modify the bond mode of a default distributed switch.

6. Click **OK**.

11.1.1.2.2 Manage Joined Hosts

Procedure

1. In the navigation pane, choose **Inventory > Network Resource**.
2. In the resource tree, select a target distributed switch.

3. On the target distributed switch's details page, click the **Uplink** tab.
4. In the **Joined Host** list, select a target host.
 - a) (Optional) To add new ports to this bond, click **Actions > Add Physical Network Interface**.

**Note:**

1. You can only add ports that are not currently bonded.
2. New ports must have the same speed as existing ports in the bond.
3. A bond supports a maximum of 8 physical ports.

- b) (Optional) To remove physical ports from the bond, click **Actions > Remove Physical Network Interface**.

**Note:**

You must keep at least 1 physical port.

- c) (Optional) To remove the host's entire uplink, click **Actions > Disconnect Host Uplink**.

**Note:**

1. Disconnecting host uplink will delete all bonds on this host. Proceed with caution.
2. This action will also detach all VM NICs on this host. Proceed with caution.
3. The default distributed switch does not support disconnecting host uplink.

11.1.1.2.3 Configure Uplinks for Hosts

Procedure

1. In the navigation pane, choose **Inventory > Network Resource**.
2. In the resource tree, select a target distributed switch.
3. On the target distributed switch's details page, click the **Uplink** tab.
4. In the **Unjoined Host** list, select a target host.
 - a) (Optional) If the new host's bond configuration matches the distributed switch's uplink configuration associated with its cluster, the host will automatically join the uplink.
 - b) (Optional) If the configuration does not match, click **Actions > Join Uplink**.
After joining, the host's bond configuration will adjust to match the switch's uplink configuration.

11.1.1.3 Network Topology

11.1.1.3.1 View Network Topology

Procedure

1. In the navigation pane, choose **Inventory > Network Resource**.
2. In the resource tree, select a target distributed switch.
3. On the target distributed switch's details page, click the **Network Topology** tab.
4. View the network topology.

The network topology displays the network relationship structure centered around the distributed switch, showing its associations with clusters, hosts, distributed port groups, and virtual machines.

11.1.1.3.2 Supported Network Topology Operations

The following table lists the actions that you can perform on the network topology.

Operation	Description
Refresh	Displays the current latest network topology.
Zoom In / Zoom Out	Zooms in or out to view 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	Views the network topology in full screen.
Highlight Display	Selects a resource and highlights the resource and its associated resources.
Hover Display	Displays relevant information about a resource when the mouse hovers over it, with support for navigating to the resource details.
Search	Searches for topology resources by resource name or UUID.

11.1.1.4 Attach/Detach a Distributed Switch to/from a Cluster

Prerequisites

The default distributed switch only supports attaching/detaching bare metal clusters.

Procedure

1. In the navigation pane, choose **Inventory > Network Resource**.

2. In the resource tree, select a target distributed switch.
3. On the target distributed switch's details page, click the **Cluster** tab.
4. (Optional) To provide network for more clusters and hosts in the data center, click **Attach Cluster**.
 - a) In the **Attach Cluster** dialog, select a target cluster, bond configuration, and host NIC.

When attaching a cluster, you cannot specify bond mode. The bond mode inherits the distributed switch's configuration directly.
 - b) Click **OK**.
5. (Optional) To detach clusters, click **Detach Cluster**.
 - a) In the **Detach Cluster** dialog, select target clusters.
 - b) Click **OK**.

**Note:**

After detaching a cluster, the corresponding VM NICs will be removed. Proceed with caution.

11.1.1.5 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 distributed port groups in the distributed switch are referenced by a VM memory snapshot.
- You cannot delete a distributed switch if distributed port groups in the distributed switch are associated with Kernel adapters.

Procedure

1. In the navigation pane, choose **Inventory > Network Resource**.
2. In the resource tree, select a target distributed switch.
3. On the target distributed switch's details page, click **Actions > Delete**.

**Note:**

Deleting a distributed switch deletes all its distributed port groups and detaches associated VM NICs. Proceed with caution.

11.1.2 Distributed Port Group

Distributed port groups are used to provide network connections to virtual machines and for Kernel adapter traffic.

After adding the first host to a cluster, NexaSphere automatically creates a default distributed switch, default distributed port group, and default Kernel adapter based on the configuration of this host, for centralized management of the host's management network. Based on your network planning, you can flexibly reuse the default distributed switch to create new distributed port groups, or use custom distributed switches to create new distributed port groups.

You can understand the supported basic operations for distributed port groups from the perspective of adding, deleting, modifying, and querying.

- [Create a Distributed Port Group](#)
- [Edit Distributed Port Group](#)
- [View Distributed Port Group](#)
- [Delete a Distributed Port Group](#)

Create a Distributed Port Group

The platform provides multiple entry points for creating distributed port groups. You can create one or more distributed port groups from the following two main entry points:

- Navigate to **Inventory > Network Resource**, right-click the target distributed switch, and then select **New Distributed Port Group**.
- Navigate to **Inventory > Network Resource**, select the target distributed switch. Then, on the right side of the platform page, click **Actions > New Distributed Port Group**, or in the **Distributed Port Group** subpage, click **New Distributed Port Group**.

To create a distributed port group, you need to complete the following information configuration:

- **Name:** The name of the distributed port group.
- **Description:** The description of the distributed port group.
- **Distributed Switch:** Select the distributed switch corresponding to the distributed port group.
- **VLAN Type:** The VLAN type of the distributed port group, choose either **None** or **Standard VLAN**. If you select Standard VLAN, you can configure the VLAN ID for the distributed port group.
- **DHCP Service:** A service for automatically assigning IP addresses to internal resources within the platform. It is disabled by default. This service only affects resources within the platform

and will not conflict with any existing DHCP servers. If enabled, you need to configure the network segment, IP allocation policy, and DHCP service IP.

- **Network Segment Method:** Supports IP range and CIDR as two types of network segments.
 - If you choose IP range, you need to enter the start IP, end IP, subnet mask, and gateway.



Note:

Do not include the gateway, broadcast address, or network address in the added IP range.

- If you choose CIDR, you need to enter the CIDR block and gateway. The gateway can be the first or last address in the CIDR block. If left blank, the first address in the CIDR block is used as the gateway by default.
- **IP Allocation Policy:** Supports random, sequential, and circular allocation policies:
 - Random Allocation: The system randomly assigns IP addresses within the network segment.
 - Sequential Allocation: The system allocates all available IP addresses within the network segment in ascending order. IP addresses released during the process will be prioritized for allocation in the next round.
 - Circular Allocation: The system allocates IP addresses within the network segment in ascending order. IP addresses released during the process will be allocated after all currently available IP addresses have been assigned once.
- **DHCP Service IP:** The IP address used by the DHCP service. The DHCP service uses this IP to assign IP addresses to resources using this distributed port group.
 - When creating a distributed port group for the first time or adding the first network segment to the distributed port group, you can customize the DHCP service IP. If a DHCP service IP already exists for the distributed port group, you cannot customize the DHCP service IP when adding a network segment.
 - The DHCP service IP must be within the CIDR of the added IP range and not be in use.
- **DNS:** DNS resolution service for the distributed port group.

After clicking **OK**, the creation is complete.

Edit Distributed Port Group

If you only need to modify the name and description of the distributed port group, you can do so on the target distributed port group page by clicking **Actions > Edit Name and Description**.

If you need to modify the MTU or IP allocation policy of the distributed port group, you can do so on the target distributed port group page by clicking **Actions > Edit Configuration**.

If you need to add more IPv4 network segments to or delete IPv4 network segment configurations from the distributed port group, you can do so respectively on the target distributed port group's **Network Segments** page.



Note:

- Deleting a network segment will cause the network cards of virtual machines using that segment to be unloaded. Proceed with caution.
- If the distributed port group already has a DHCP service IP and that IP is within the selected network segment, deleting all network segments under the distributed port group will release the DHCP service IP. Deleting only some network segments will leave the DHCP service IP unchanged.

If you need to add more DNS configurations to or delete DNS configurations from the distributed port group, you can do so respectively on the target distributed port group's **DNS** page.

View Distributed Port Group

If you want to obtain usage rate data for IP addresses of NICs in a distributed port group over various time periods, you can view this information on the target distributed port group's **Monitoring** page:

- This page displays a line chart showing the percentage of used and available IP addresses within the selected time range. The chart updates in real-time, facilitating analysis of network resources.
- You can select the time range as needed, including: 15 minutes, 1 hour, 6 hours, 1 day, 1 week, 1 month, 1 year, and a custom time range.

Delete a Distributed Port Group

If you have determined that you no longer need an existing distributed port group, you can delete it on the target distributed port group page by clicking **Actions > Delete**. You can also delete distributed port groups in bulk on the data center resource **Network > Network Resource >**

Distributed Port Group page or on the distributed switch resource **Distributed Port Group** page.



Note:

Deleting a distributed port group will unload the network cards of virtual machines using this port group. Proceed with caution.

11.1.2.1 Create a Distributed Port Group

Procedure

1. In the navigation pane, choose **Inventory > Network Resource**.
2. In the resource tree, right-click a target distributed switch and select **New Distributed Port Group**.
3. In the **New Distributed Port Group** dialog, set the following parameters:
 - **Name:** Enter a name for the distributed port group.
 - **Description:** Enter a brief description for the distributed port group.
 - **Distributed Switch:** Select a distributed switch for the distributed port group.
 - **VLAN Type:** Select a VLAN type. When selecting Standard VLAN, you need to specify a VLAN ID.
 - **DHCP Service:** Choose whether to enable the automatic IP address assignment for platform resources.
 - **IP Address Management:** Choose whether to enable the IP address management. When enabled, you can add network ranges to this distributed port group. IP addresses in these ranges can be allocated via DHCP service (when enabled) to resources in the network.
 - **IP Address Type:** Supports IPv4 and IPv6.

When selecting IPv4, configure the following parameters:

- **IP Allocation Policy:** After enabling the DHCP service, you can select one of the following policies to assign IP addresses.
 - **Random Allocation:** The system randomly assigns IP addresses from the network range.
 - **Allocate in Order:** The system assigns all available IP addresses from the 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 available IP addresses from the network range in ascending order. Released IP addresses are assigned when currently available IP addresses are used up.

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:** Supports IP range and CIDR.
 - For IP range, enter start IP, end IP, netmask, and gateway.



Note:

Do not include gateway, broadcast address, or network addresses in the IP range.

- For CIDR, enter CIDR Block and gateway. You can enter the first or last CIDR address for gateway. If left blank, the first CIDR address will be used.

When selecting IPv6, configure the following parameters:

- **Network Range Method:** Supports IP range and CIDR.
 - For IP range, enter start IP, end IP, prefix length, and gateway.
 - For CIDR, enter CIDR Block.
- **IP Configuration Mode:** Select IPv6 address allocation method.
 - **Stateful-DHCP (Default):** The interface address and other parameters are all configured through DHCP. The IP range method supports stateful DHCP.
 - **Stateless-DHCP:** The interface address is automatically derived from the route advertisement prefix and the interface Mac address. Other parameters are configured through DHCP.
 - **SLAAC:** The interface address is automatically derived from the prefix of the route advertisement that also contains other parameters.
- **DHCP IP:** Specify an IP address used by the DHCP service.
- **DNS:** Specify a DNS to provide DNS resolution service for the distributed port group.

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

11.1.2.2 Modify Distributed Port Group Configurations

Procedure

1. In the navigation pane, choose **Inventory > Network Resource**.
2. In the resource tree, select a target distributed port group.
3. (Optional) To modify basic configurations, click **Actions > Modify Configuration**.
 - a) In the **Modify Configuration** dialog, modify the name, description, VLAN ID, MTU, DHCP service as needed.
 - b) Click **OK**.
4. (Optional) To add or delete network ranges, click the **Network Range** tab and perform the corresponding operations as needed.



Note:

- Deleting a network range also detaches the VM NICs that are using the network range.
- If the distributed port group has a DHCP IP that is in the selected network range, deleting the network range does not delete the DHCP IP. If all network ranges under the distributed port group are deleted, the DHCP IP will be released.

5. (Optional) To modify the DNS, click the **DNS** tab and perform the corresponding operations as needed.

11.1.2.3 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 a target distributed port group.
3. On the distributed port group's details page, click **Actions > Delete**.



Note:

Deleting a distributed port group detaches the VM NICs that are using this network. Proceed with caution.

4. After acknowledging the risk, click **OK**.

11.1.3 Kernel Adapter

The Kernel adapter uses network labels to identify physical network traffic. After successfully adding a host to a cluster, NexaSphere creates a default Kernel adapter for each host, which is used to obtain and display the management network of that host. You can create Kernel adapters for hosts to manage storage network traffics.

11.1.3.1 Create a Kernel Adapter

Procedure

1. In the navigation pane, choose **Inventory > VM and Host**.
2. In the resource tree, select a target host.
3. On the details page of the target host, click **Kernel Adapter**.
4. On the **Kernel Adapter** tab, click **New Kernel Adapter**.
5. In the **New Kernel Adapter** dialog, set the following parameters:
 - **Name**: Enter a name for the Kernel adapter.
 - **Description**: Enter a brief description for the Kernel adapter.
 - **Network Service**: Display **Storage** by default, indicating this Kernel adapter handles storage network traffics.
 - **Distributed Port Group**: Select a distributed port group.
 - **IPv4 Address**: Assign an IPv4 address for the Kernel adapter.
 - **Netmask**: Specify a netmask.
6. Review the configuration and click **OK**.

11.1.3.2 Modify Kernel Adapter Configurations

Procedure

1. In the navigation pane, choose **Inventory > VM and Host**.
2. In the resource tree, select a target host.
3. On the details page of the target host, click **Kernel Adapter**.
4. On the **Kernel Adapter** tab, choose a Kernel adapter from the list and click **Actions > Modify Configuration**.
5. In the **Modify Configuration** dialog, perform the corresponding modifications as needed.

- For the default Kernel adapter, you can only modify the name, description, and whether to select **Storage** network service. When 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**.

11.1.3.3 Delete a Kernel Adapter

Procedure

1. In the navigation pane, choose **Inventory > VM and Host**.
2. In the resource tree, select a target host.
3. On the details page of the target host, click **Kernel Adapter**.
4. On the **Kernel Adapter** tab, choose a Kernel adapter from the list and click **Actions > Delete**.



Note:

- You cannot delete a default Kernel adapter.
- Deleting Kernel adapters will release the associated IP addresses. Storage services dependent on these IPs will be interrupted. Proceed with caution.

5. After acknowledging the risks, click **OK**.

11.2 Network Service

NexaSphere provides security group network services to ensure the security of east-west traffic between virtual machines.

11.2.1 Security Group

This section describes how security groups work and how to use them:

- [Overview](#)
- [Create a Security Group and Related Rules](#)
- [Modify Security Groups and Related Rules](#)

11.2.1.1 Overview

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.

Functional Framework

Security groups control traffic to and from network interfaces through security rules within the group. A network interface can be part of multiple security groups, and by setting the priority of security groups, traffic is first matched against rules in the higher-priority groups.

A security group can contain multiple security rules. Based on their creation mechanism, these can be divided into system rules and custom rules:

- **System Rules:** After creating a new security group, the system provides two default rules:
 - **Intra-Group Communication Rule:** Network interfaces within the same security group are allowed to communicate with each other by default. This rule has a higher priority than all custom rules and cannot be modified or deleted, only disabled.
 - **Intra-/Inter-Group Communication Rule:** Network interfaces within the security group are allowed to access interfaces outside the security group by default, but interfaces outside the group are not allowed to access those inside by default. This rule has a lower priority than all custom rules and supports modifying the default intra-group and inter-group access behavior for individual virtual machine network interfaces.
- **Custom Rules:** Rules added to the security group by users.

The security group rules consist of direction, target, action, protocol & port, and priority:

- **Rule Direction:** Security group rules primarily control the source or destination of traffic. Based on the direction of traffic flow, they can be categorized as inbound rules and outbound rules:
 - **Inbound Rule:** For traffic entering the network interface from the outside, primarily controlling the source of traffic.
 - **Outbound Rule:** For traffic sent out from the network interface, primarily controlling the destination of traffic.
- **Rule Target:** The target of the security group rule (inbound/outbound rule), including source and destination:
 - **Source:** Corresponds to the inbound rule, supporting the use of IP addresses/ranges or security groups as sources. Inbound rules allow/reject traffic from the specified IP addresses /ranges or security groups.
 - **Destination:** Corresponds to the outbound rule, supporting the use of IP addresses/ranges or security groups as destinations. Outbound rules allow/reject traffic from the current group's network interfaces to the target IP addresses/ranges or security groups.

- **Action:** The specific action taken for traffic matching the rule conditions, including Allow and Deny:

- Allow: Allows network request traffic to flow into or out of the network interface.
- Deny: Does not allow network request traffic to flow into or out of the network interface.

By default, if traffic entering or leaving the network interface does not match any custom rules, inbound traffic is denied and outbound traffic is allowed.

- **Protocol and Port:** The packet protocol and corresponding port targeted by the rule. Protocols include ALL, TCP, UDP, and ICMP:
 - ALL: Indicates coverage of all protocol types, and ports cannot be specified.
 - TCP: Supports ports 1-65535.
 - UDP: Supports ports 1-65535.
 - ICMP: Does not support specifying a port.
- **Priority:** The relative precedence of one security group rule over others, with supported values ranging from 1 to 100. Higher numbers indicate lower priority.

11.2.1.2 Create a Security Group and Related Rules

To use security groups with your virtual machine, you need to create a new security group, add rules to the security group, and bind the security group to the NIC:

1. [Create a Security Group](#)
2. [Add Rules](#)
3. [Associate VM NIC](#)

Create a Security Group

On the NexaSphere platform, select the target data center, then click **Network > Security Group**, and follow the configuration below to create a new security group:

- **Name:** Enter the name for the security group.
- **Description:** Provide a description for the security group.

After clicking **OK**, the new security group will be created.

Add Rules

After creating a new security group, you can add individual or batch inbound and outbound rules to the security group on its **Overview** page.

- [Add Individually](#)

- [Add in Batch](#)

Add Individually: Depending on the direction of the rule you wish to add, select the **Ingress Rule** or **Egress Rule** tab, and click the **Add Rule** button. Follow the example below to add the rule:

- **Type:** The direction of traffic controlled by the rule, displayed as Ingress or Egress.
- **Priority:** The priority of the rule, which automatically increments by 1 for each new rule added. Higher numbers indicate lower priority.
- **IP Address Type:** Supports IPv4 address type.
- **Protocol:** The communication protocol targeted by the rule, supporting ALL, TCP, UDP, and ICMP.
- **Port:** When selecting TCP or UDP, specify the port targeted by the rule:
 - If specifying a range of ports, use the format `Start Port - End Port`.
 - If specifying multiple ports or ranges, separate them with an English comma “,”. You can specify up to 10 ports or ranges.
- **Source:** Required when adding an inbound rule, indicating whether to allow or deny traffic from the specified IP addresses/ranges or security groups:
 - When specifying by IP address/range, you can enter a range using the format `Start IP - End IP`.
 - When specifying by IP address/range, you can enter CIDR notation. If specifying CIDR along with other types of IP addresses, the CIDR mask must be 24 bits. If specifying only CIDR, there is no limit on the mask.
 - If specifying multiple IP addresses/ranges, separate them with an English comma “,”.
- **Destination:** Required when adding an outbound rule, indicating whether to allow or deny traffic from the NICs in this group to the specified IP addresses or security groups:
 - When specifying by IP address/range, you can enter a range using the format `Start IP - End IP`.
 - When specifying by IP address/range, you can enter CIDR notation. If specifying CIDR along with other types of IP addresses, the CIDR mask must be 24 bits. If specifying only CIDR, there is no limit on the mask.
 - If specifying multiple IP addresses/ranges, separate them with an English comma “,”.
- **State:** Whether to enable the rule immediately after creating the security group. By default, the rule is enabled. If set to disabled, the interfaces in the group will not match this rule until you manually enable it.

- **Description:** Description of the security group rule.

After clicking **OK**, the rule will be successfully added.

Batch Add: You can add multiple inbound and outbound rules to a security group by importing rules:

- Click **Actions > Import Rule**, upload a CSV file, and click **OK**. The rules will be imported successfully.



Note:

- The imported rules will not affect existing rules. Their priority will default to be placed after existing rules and will be in a disabled state.
- After import, users can manually adjust the priority and enable these rules.
- To ensure system compatibility, the imported file must be edited using Microsoft Excel.
- If you need to reuse the rules of one security group for another security group, you can do the following on the security group page, depending on the scenario:
 - If you only need to export inbound rules or outbound rules, you can click the download button at the top right of the rule list under the corresponding tab and choose to export the current page or all . This will export the rules in CSV format.
 - If you need to export all inbound and outbound rules, click **Actions > Export Rule**. This will export the rules in CSV format.

After exporting, you can import them into the target security group following the steps above.

Associate VM NIC

After adding rules to a security group, you can choose the following paths to bind the security group to virtual machine NICs based on your scenario:

- If you need to bind the security group to multiple virtual machine NICs, you can select target interfaces in bulk on the **VM NIC** page of the security group, and follow the example below to bind them:
 - **Network:** Select the network scope applicable to the security group, choosing either all distributed port groups or specific port groups in the data center.
 - **NIC:** Select target NIC to bind.



Note:

If the IP address of a virtual machine NIC is shown as empty on the platform, the default security group rule (intra-group communication rule) will not apply to that interface.

- If you need to bind multiple security groups to a virtual machine NIC, you can do so on the **Overview** page of the target virtual machine, using **Edit Configuration** to bind security groups to the virtual machine NICs in bulk. The smaller the number assigned to a security group, the higher its priority for taking effect.



Note:

Please configure carefully to avoid conflicts between rules across different security groups.

11.2.1.3 Modify Security Groups and Related Rules

After creating a new security group, its rules, and binding them to virtual machine network interfaces, if you need to modify or delete the security group and related rules, refer to the content in this section for guidance.

- [Security Group](#)
- [Security Group Rules](#)

Security Group

If you need to change the enabled status of a security group, you can use the **Enable** and **Disable** operations to make the changes.

If you need to modify the name and description of a security group, you can click **Edit Name and Description** in the operation column of the target security group to make the changes.

If you need to unbind virtual machine network interfaces from a security group, follow the instructions below based on your scenario:

- If you need to unbind multiple virtual machine network interfaces from a security group, you can select the target interfaces on the **VM NIC** page of the security group and click **Disassociate VM NIC** to unbind them.
- If you need to unbind multiple virtual machine network interfaces from a security group, you can use **Edit Configuration** on the **Overview** page of the target virtual machine to unbind the security group from the virtual machine network interfaces in bulk.

If you are certain that you no longer need an existing security group, you can click **Delete** in the operation column of the target security group to delete it.

**Note:**

Deleting a security group will also delete any custom rules created for the security group. Please proceed with caution.

Security Group Rules

If you wish to modify the system default rules, note that security group rules are divided into system default rules and custom rules. The system default rule is the intra-group communication rule, which only supports the disable operation.

If you need to perform operations on custom-added or imported security group rules, you can go to the **Overview** page of the target security group and select the **Ingress Rule** or **Egress Rule** tab as needed, and perform the following operations based on your scenario:

- If you need to modify the enabled status of individual or multiple security group rules, you can use the **Enable** and **Disable** operations to make the changes.
- If you need to modify the configuration of a single security group rule, such as priority, allow or deny policy, protocol and port, source/destination, enabled status, and description, you can click **Edit Rule** in the operation column of the target security group and make the necessary changes.
- If you need to adjust the priority of individual or multiple security group rules, you can click the **Adjust Priority** button and adjust the relative priority order of the rules as needed by dragging. Value range: integers within the range of 1-100. Higher numbers indicate lower priority.
- If you are certain that you no longer need a particular security group rule or several rules, select the target rule(s) and click **Delete** to delete them.

12 O&M Management

12.1 Resource Monitoring

12.1.1 Resource Performance Monitoring

12.1.1.1 Overview

NexaSphere provides visual charts that display various monitoring data for resources over a period of time. These charts include multiple key performance monitoring metrics, helping you gain an intuitive understanding of resource performance conditions.

Monitoring Chart Types

Chart Type	Description
Bar Chart	Displays monitoring data of resource capacity load in the form of proportional bars, providing an intuitive understanding of resource capacity information.
Line Chart	Displays monitoring data of various loads on resources in the form of a line chart, offering an intuitive understanding of resource health status.

Monitoring Data Collection Methods

NexaSphere provides two monitoring methods for virtual machines. Generally speaking, for memory data, Advanced Monitoring offers better accuracy than Basic Monitoring. It is recommended to use Advanced Monitoring when monitoring memory data.

- **Basic Monitoring:** Monitoring data is obtained from the host via Libvirt.
- **Advanced Monitoring:** Monitoring data is obtained from the virtual machine by an advanced monitoring agent. VMTools must be pre-installed on the virtual machine for this method.

Monitoring Data Collection Intervals

NexaSphere uses real-time monitoring, with resource monitoring charts refreshing data every 10 seconds by default.

12.1.1.2 Capacity Monitoring

NexaSphere provides information on the usage and allocation of various computing and storage resources, including virtual machines, hosts, clusters, data storage, data centers, and root nodes

(management nodes). This allows you to comprehensively understand the platform's resource usage from both micro and macro perspectives.

Capacity Monitoring Metrics

You can go to the overview details page of the corresponding resource to understand the platform's resource usage from the **Capacity Information** card. The following table lists the detailed monitoring metrics for various resources.

Object	Monitoring Metrics and Description
Root Node	• CPU: Total physical CPU GHz and average utilization rate across all data centers. • Memory: Total physical memory, average utilization rate, and remaining available capacity across all data centers. • Storage: Total physical storage, average utilization rate, and remaining available capacity across all data centers.
Data Center	• CPU: Total physical CPU GHz and average utilization rate within the data center. • Memory: Total physical memory, average utilization rate, and remaining available capacity within the data center. • Storage: Total physical storage, average utilization rate, and remaining available capacity within the data center.
Data Storage	• Storage Utilization: Total storage resources, utilization rate, and remaining available capacity. • Storage Allocation Ratio: Allocation status of storage resources. • Storage Distribution: Distribution of storage resources, including: total capacity after overcommitted, reserved capacity, allocated capacity (such as snapshot capacity, image cache, migration storage, virtual machine disk capacity.), and remaining allocatable capacity.
Cluster	• Resource Utilization: Total physical CPU and memory resources, utilization rate, and remaining available capacity in the cluster. • Resource Allocation Ratio: Allocation status of physical CPU and memory resources in the cluster. • Resource Distribution: Distribution of CPU and memory resources after overcommitted in the cluster, including: total capacity after overcommitted, reserved capacity, allocated capacity, and remaining allocatable capacity.

Object	Monitoring Metrics and Description
Host	• Resource Utilization: Total physical CPU, memory, and storage resources on the host, utilization rate, and remaining available capacity. • Resource Allocation Ratio: Allocation status of CPU, memory, and storage resources on the host. • Resource Distribution: Distribution of CPU, memory, and storage resources after overcommitted on the host, including: total capacity after overcommitted, reserved capacity, allocated capacity, and remaining allocatable capacity.
Virtual Machine	• CPU: Number of CPU cores and utilization rate for the virtual machine. • Memory: Total memory capacity, used capacity, and remaining available capacity for the virtual machine. • Storage: Total storage capacity, used capacity, and remaining available capacity for the virtual machine.

Capacity Calculation Rules

Category	Calculation Rules
Resource Utilization Rate	Total CPU = Physical Cores × Single-Core GHz
Resource Allocation Ratio	• Allocation Ratio = Allocated : Total Overcommit Capacity • Total Overcommit Capacity = Physical Total – Reserved Physical Capacity • Total Allocatable = Total Overcommit Capacity × Overcommit Ratio • Free to Allocate = Total Allocatable – Allocated
Resource Distribution	CPU • CPU Overcommitted Total = Physical CPU Total × Overcommit Ratio Memory • Memory Overcommitted Total = Reserved Memory + Total Allocatable Memory Capacity • Total Allocatable Memory = (Physical Memory Total – Reserved Memory) × Overcommit Ratio Storage • Storage Overcommitted Total = Reserved Capacity + Total Allocatable Storage Capacity

Category	Calculation Rules
	Total Allocatable Storage = (Physical Storage Total – Reserved Capacity) × Overcommit Ratio

The meaning of overcommitment and allocation are as follows:

- **CPU Overcommitment:** This indicates that a single physical CPU core can be virtually divided into N logical CPU cores for allocation to virtual machines.

For example, if the CPU overcommitment ratio is 2:1, then one physical CPU core can be virtually divided into 2 logical CPU cores. Therefore, if a host has 10 physical CPU cores, it can be virtually divided into 20 logical CPU cores for allocation to virtual machines.

- **Memory/Storage Overcommitment:** This indicates that a unit of memory/storage capacity can be virtually expanded into N units of memory/storage capacity for allocation to virtual machines.

For example, if the memory/storage overcommitment ratio is 2:1, then 1 GB of memory/storage capacity can be virtually expanded into 2 GB. Therefore, if a host has 100 GB of memory/storage, it can be virtually expanded into 200 GB for allocation to virtual machines.

- **CPU Allocation:** This indicates that a physical CPU core is actually virtually divided into N logical CPU cores for use by virtual machines. Therefore, the CPU allocation ratio $\leq$ CPU overcommitment ratio.

For example, if the CPU allocation ratio is 1.5:1, then one physical CPU core is actually virtually divided into 1.5 logical CPU cores. Therefore, if a host has 10 physical CPU cores, they have actually been virtually divided into 15 logical CPU cores for allocation to virtual machines.

- **Memory/Storage Allocation:** This indicates that a unit of memory/storage capacity is actually virtually expanded into N units of memory/storage capacity. Therefore, the memory/storage allocation ratio $\leq$ memory/storage overcommitment ratio.

For example, if the memory/storage overcommitment ratio is 1.5:1, then 1 GB of memory/storage capacity is actually virtually expanded into 1.5 GB. Therefore, if a host has 100 GB of memory/storage, it has actually been virtually expanded into 150 GB for allocation to virtual machines.

Using host storage as an example, if the total physical storage capacity is 100 GB, the reserved physical capacity is 10 GB, the overcommitment 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 × 2 = 180 GB
- Remaining Allocatable Storage = 180 GB - 150 GB = 30 GB

12.1.1.3 View Monitoring Charts

NexaSphere supports visualizing load monitoring data for various resources in the form of line charts. This not only helps you quickly understand the inventory of computing, storage, and network resources for resource objects but also provides an intuitive understanding of resource health conditions.

Procedure

1. In the navigation pane, choose **Inventory**.
2. Select a valid resource object, such as a virtual machine, host, cluster, image storage, data storage, or distributed port group.
3. In the right-side pane, click **Monitoring**.
4. (Optional) Select the monitoring items you want to display.
5. (Optional) Choose or customize the time range.
6. (Optional) Select one or multiple monitoring objects.

12.1.1.4 Customize Monitoring Charts

You can customize monitoring charts to view more monitoring data.

- **Details:** Hover the mouse over the chart to display detailed information about the relevant data points.
- **Custom Time Span:** By default, it displays monitoring data for the past 15 minutes. Valid values include 15 minutes, 1 hour, 6 hours, 1 day, 1 week, 1 month, 1 year, and custom.
- **Custom Monitoring Items:** Flexibly select the monitoring metrics you want to focus on based on your business needs.
- **Custom Monitoring Objects:** Display data for all or specified monitoring objects.
- **Custom Chart Position:** Freely drag and rearrange the position of monitoring charts.

Appendix of Monitoring Items

Object	Metric	Item and Description
Cluster	• CPU • Memory	• CPU Utilization Sum • Memory Usage Percentage

Object	Metric	Item and Description
	- Disk - NIC	- Disk IOPS Sum - NIC Data Transfer Rate Sum
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, 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, 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.
	Memory	Memory usage: The amount of used and free resource memory.
	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 the platform system files. Disk Usage of Platform System Files: The amount of disk capacity occupied by the platform system files.
	NIC	NIC Data Transfer Rate: The current send and receive rate of the resource's NIC.

Object	Metric	Item and Description
		• NIC Packet Rate: The current send and receive packet rate of the resource's NIC. • NIC Packet Discard Rate: The current packet drop rate for outgoing and incoming packets on the resource's NIC.
Virtual Machine	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, 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, 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.
	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. • Memory Utilization: The percentage of resource memory that is currently in use.
	Disk	• Disk Speed: The read and write speed of the resource disk. • Disk IOPS: The read and write IOPS of the resource disk.

Object	Metric	Item and Description
		• 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.
	NIC	• NIC Data Transfer Rate: The current send and receive rate of the resource's NIC. • NIC Packet Rate: The current send and receive packet rate of the resource's NIC. • NIC Packet Discard Rate: The current packet drop rate for outgoing and incoming packets on the resource's NIC.
Data Storage	Capacity	Capacity Percent Used: The percentage of capacity currently used by the resource.
Image Storage - Standalone Image Storage/ Distributed Image Storage	Capacity	Capacity Percent Used: The percentage of capacity currently used by the resource.
Image Storage - Standalone Image Storage	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, 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, 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

Object	Metric	Item and Description
		hard disk drive to load data into memory after initiating a read or write operation.
	Disk	• Disk Speed: The read and write speed of the resource disk. • Disk IOPS: The read and write IOPS of the resource disk.
	Memory	Memory Usage: The amount of used and free resource memory.
	NIC	• NIC Data Transfer Rate: The current send and receive rate of the resource's NIC. • NIC Packet Rate: The current send and receive packet rate of the resource's NIC. • NIC Packet Discard Rate: The current packet drop rate for outgoing and incoming packets on the resource's 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.

12.1.2 Dashboard Monitoring

NexaSphere The dashboard displays platform resource status statistics, platform load trends, platform usage statistics, resource top rankings, user information, and unread alarm statistics for the past seven days in a card format.

- Each time you enter or refresh the dashboard, the latest data is fetched and displayed in real-time. Additionally, chart-based modules automatically refresh data every 30 seconds by default.
- The dashboard by default shows the resource data for the current data center. You can click the switch button in the top left corner of the page to specify which data center's resource data to display.
- Status charts use a standardized color scheme: green indicates normal status, red indicates an abnormal status, and gray indicates other statuses.
- Percentage progress bars are color-coded as blue (less than 60%), yellow (greater than or equal to 60% but less than 80%), and red (greater than or equal to 80%) to visually represent the current resource usage state.

- For resource status cards and some load trend and usage statistics cards, you can click on the resource name or statistical numbers to navigate to the corresponding resource page.

12.1.3 Dual Management Node Monitoring

If your environment consists of two management nodes, navigate to **Reliability > MN Monitoring** page to view the management node monitoring data.

Before you check the management node monitoring data, you should be aware of the following information:

- This page uses three colors: green, red, and gray. Green indicates normal status, while other colors indicate abnormal status.
- The dual-management node setup follows a active-standby model, with only one active management node. The node displaying VIP is the active management node, and the one without VIP is the standby management node.
- If the standby management node is in an abnormal state, the active management node will fail to switch and the management nodes will go down. Therefore, address any management node issues promptly.

The management node monitoring displays the management node IPs, node status, VIP, and management service status for multiple management nodes. The main services monitored include the following:

- **Arbiter Gateway Reachable:**

Monitors whether the arbitration IP of the active-standby management node is reachable. If unreachable, it may cause the high availability feature of the management node to fail.

- **Peer MN Reachable:**

Monitors whether the standby management node is reachable. If the standby management node is unreachable, communication with the standby node will not be possible.

- **VIP Reachable:**

Monitors whether the VIP is reachable. If the VIP is unreachable, the primary management node cannot access the UI interface via the VIP.

- **Database Status:**

Monitors the status of the database. If the database is abnormal, there may be a risk of data loss. Please restore the fault promptly.

12.1.4 Host Hardware Monitoring

NexaSphere supports monitoring the status of host hardware components such as CPU, memory, sensors, PCIe devices, and more.

The hardware components that can be monitored on the host include:

- CPU
- Memory
- Physical Disks
- Physical Network Cards
- GPU Devices
- Block Devices
- USB Devices
- Sensors (Voltage, Current, Fans, Temperature)
- Power Supply
- PCIe Devices

12.2 Alarm Service

12.2.1 Overview

NexaSphere supports setting up alarms for resources based on load and capacity usage, as well as configuring event alarms for predefined events occurring within the platform. When critical resources experience abnormalities, the platform instantly pushes alarm messages to designated endpoints to help quickly identify and resolve issues, minimizing potential business disruptions.

Alarm Service Infrastructure

The NexaSphere alarm service consists of two main components: the monitoring system and the notification service.

- **Monitoring System**
 - Provides time-series data monitoring and alarms, including resource load and capacity data.
 - Captures predefined events from the platform and triggers alarms.
 - Supports custom alarms and alarm message templates.
 - Allows viewing of alarm messages through multiple entry points.
- **Notification Service**

- Sends alarm notifications to designated endpoints, such as system, email, DingTalk, Lark, WeCom, HTTP applications, Microsoft Teams, and SNMP Trap receivers.

Usage Recommendations

Considering that monitoring data consumes certain system resources, it is recommended to configure NexaSphere related resources according to the following requirements:

- Plan a dedicated physical server as the management node for the platform.
- Given that monitoring data may periodically consume system disk I/O resources, it is recommended to use an SSD for the management node's system disk.
- To avoid excessive system disk usage due to large monitoring data, plan for a system disk space of at least **1TB**.
- If your system disk space is less than 500GB, you can modify the following configurations in **System Parameters**:
 - **Monitoring Data Retention Period**: Set to 1 month.
 - **Monitoring Data Retention Size**: Set to a power of 2, such as 32GB, 64GB, or 128GB.

12.2.2 Endpoint

12.2.2.1 New Endpoint

Endpoints are the foundation for the notification service to push alarm messages. NexaSphere provides a system alarm endpoint as the default option, and you can also create custom endpoints of various types.

- [Email](#)
- [DingTalk](#)
- [Lark](#)
- [WeCom](#)
- [SMS](#)
- [HTTP Application](#)
- [Microsoft Teams](#)
- [SNMP Trap Receiver](#)

12.2.2.1.1 Email

Prerequisites

An email server needs to be added in advance. For more information, see [Add Email Server](#).

Procedure

1. In the navigation pane, choose **O&M Management > Endpoint**.
2. On the **Endpoint** page, click **New Endpoint**.
3. In the **New Endpoint** dialog, set the following parameters:
 - **Name**: Set the name for the endpoint.
 - **Description**: Optionally fill in a description for the endpoint.
 - **Type**: Select **Email**.
 - **Email Server**: Choose an email server that has been added.
 - **Email Address**: Enter the email address.
 - **Message Language**: Select the notification language for alarm messages, including Simplified Chinese, English.
4. Review the configuration and click **OK**.

What's next

- You can proceed to set up alarm message templates to ensure that alarm messages are sent out in a uniform format as specified. For more information, see [Message Template](#).
- You can proceed to create a new alarm to push resource alarm messages to designated endpoints. For more information, see [Alarm](#).

12.2.2.1.2 DingTalk

Prerequisites

- Make sure NexaSphere has an installed Advanced Edition license, and that the license is in a valid state.
- Add a DingTalk group bot in advance and configure security settings as needed. After adding , obtain the bot's Webhook URL. For more information, refer to the official DingTalk Open Platform documentation.

Procedure

1. In the navigation pane, choose **O&M Management > Endpoint**.
2. On the **Endpoint** page, click **New Endpoint**.
3. In the **New Endpoint** dialog, set the following parameters:
 - **Name**: Set the name for the endpoint.
 - **Description**: Optionally fill in a description for the endpoint.

- **Type:** Select **DingTalk**.
 - **Address:** Enter the Webhook URL of the DingTalk bot.
 - **Security Setting:** Choose the security settings configured for the DingTalk group 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:** Set whether to notify specific members when alarm messages are pushed to the DingTalk group. Options include Nobody, All, or Specified Members. When choosing Specified Members, add the mobile phone numbers of the group members.
 - **Message Language:** Select the notification language for alarm messages, including Simplified Chinese and English.
4. Review the configuration and click **OK**.

What's next

- You can proceed to set up alarm message templates to ensure that alarm messages are sent out in a uniform format as specified. For more information, see [Message Template](#).
- You can proceed to create a new alarm to push resource alarm messages to designated endpoints. For more information, see [Alarm](#).

12.2.2.1.3 Lark

Prerequisites

- Make sure NexaSphere has an installed Advanced Edition license, and that the license is in a valid state.
- Add a Lark group bot in advance and configure security settings as needed. After adding, obtain the bot's Webhook URL. For more information, refer to the official Lark Open Platform documentation.

Procedure

1. In the navigation pane, choose **O&M Management > Endpoint**.

2. On the **Endpoint** page, click **New Endpoint**.

3. In the **New Endpoint** dialog, set the following parameters:

- **Name:** Set the name for the endpoint.
- **Description:** Optionally fill in a description for the endpoint.
- **Type:** Select **Lark**.
- **Address:** Enter the Webhook URL of the Lark bot.
- **Security Setting:** Choose the security settings configured for the Lark group 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 Members:** Set whether to notify specific members when sending alarm messages via the bot. Options include Nobody, All, or Specified Members. When choosing Specified Members, add the user IDs of the designated members.
- **Message Language:** Select the notification language for alarm messages, including Simplified Chinese and English.

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

What's next

- You can proceed to set up alarm message templates to ensure that alarm messages are sent out in a uniform format as specified. For more information, see [Message Template](#).
- You can proceed to create a new alarm to push resource alarm messages to designated endpoints. For more information, see [Alarm](#).

12.2.2.1.4 WeCom

Prerequisites

- Make sure NexaSphere has an installed Advanced Edition license, and that the license is in a valid state.

- Add a WeCom group bot in advance. After adding, obtain the bot's Webhook URL. For more information, refer to the official WeCom documentation.

Procedure

1. In the navigation pane, choose **O&M Management > Endpoint**.
2. On the **Endpoint** page, click **New Endpoint**.
3. In the **New Endpoint** dialog, set the following parameters:
 - **Name**: Set the name for the endpoint.
 - **Description**: Optionally fill in a description for the endpoint.
 - **Type**: Select **WeCom**.
 - **Address**: Enter the Webhook URL of the WeCom bot.
 - **Mention Members**: Set whether to notify specific members when sending alarm messages via the bot. Options include Nobody, All, or Specified Members. When choosing Specified Members, add the user IDs of the designated members.
 - **Message Language**: Select the notification language for alarm messages, including Simplified Chinese and English.
4. Review the configuration and click **OK**.

What's next

- You can proceed to set up alarm message templates to ensure that alarm messages are sent out in a uniform format as specified. For more information, see [Message Template](#).
- You can proceed to create a new alarm to push resource alarm messages to designated endpoints. For more information, see [Alarm](#).

12.2.2.1.5 SMS

Prerequisites

Obtain an AccessKey that includes SMS services from a third-party provider in advance.

Procedure

1. In the navigation pane, choose **O&M Management > Endpoint**.
2. On the **Endpoint** page, click **New Endpoint**.
3. In the **New Endpoint** dialog, set the following parameters:
 - **Name**: Set the name for the endpoint.
 - **Description**: Optionally fill in a description for the endpoint.

- **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 phone number to receive SMS messages.

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

What's next

- You can proceed to set up alarm message templates to ensure that alarm messages are sent out in a uniform format as specified. For more information, see [Message Template](#).
- You can proceed to create a new alarm to push resource alarm messages to designated endpoints. For more information, see [Alarm](#).

12.2.2.1.6 HTTP Application

Prerequisites

Prepare the Webhook URL for the HTTP application in advance, and configure the username and password as needed.

Procedure

1. In the navigation pane, choose **O&M Management > Endpoint**.
2. On the **Endpoint** page, click **New Endpoint**.
3. In the **New Endpoint** dialog, set the following parameters:
 - **Name:** Set the name for the endpoint.
 - **Description:** Optionally fill in a description for the endpoint.
 - **Type:** Select **HTTP Application**.
 - **Address:** Enter the URL of the HTTP service.
 - **Username:** Enter the username configured for the HTTP application.
 - **Password:** Enter the password corresponding to the username.
4. Review the configuration and click **OK**.

What's next

- You can proceed to set up alarm message templates to ensure that alarm messages are sent out in a uniform format as specified. For more information, see [Message Template](#).
- You can proceed to create a new alarm to push resource alarm messages to designated endpoints. For more information, see [Alarm](#).

12.2.2.1.7 Microsoft Teams

Prerequisites

- Make sure NexaSphere has an installed Advanced Edition license, and that the license is in a valid state.
- Add the Incoming Webhook app to Microsoft Teams in advance. After adding, obtain the Webhook URL. For more information, refer to the official Microsoft Teams documentation.

Procedure

1. In the navigation pane, choose **O&M Management > Endpoint**.
2. On the **Endpoint** page, click **New Endpoint**.
3. In the **New Endpoint** dialog, set the following parameters:
 - **Name**: Set the name for the endpoint.
 - **Description**: Optionally fill in a description for the endpoint.
 - **Type**: Select **Microsoft Teams**.
 - **Address**: Enter the Webhook URL obtained from Microsoft Teams.
 - **Message Language**: Select the notification language for alarm messages, including Simplified Chinese and English.
4. Review the configuration and click **OK**.

What's next

- You can proceed to set up alarm message templates to ensure that alarm messages are sent out in a uniform format as specified. For more information, see [Message Template](#).
- You can proceed to create a new alarm to push resource alarm messages to designated endpoints. For more information, see [Alarm](#).

12.2.2.1.8 SNMP Trap Receiver

Prerequisites

- Make sure NexaSphere has an installed Advanced Edition license, and that the license is in a valid state.
- Enable SNMP management and add an SNMP Trap receiver in advance. For more information, see [Enable SNMP Management](#).

Procedure

1. In the navigation pane, choose **O&M Management > Endpoint**.

2. On the **Endpoint** page, click **New Endpoint**.
3. In the **New Endpoint** dialog, set the following parameters:
 - **Name**: Set the name for the endpoint.
 - **Description**: Optionally fill in a description for the endpoint.
 - **Type**: Select **SNMP Trap Receiver**.
 - **SNMP Trap Receiver**: Choose the SNMP Trap receiver that has been added.
4. Review the configuration and click **OK**.

12.2.2.2 Manage Endpoint

- [Modify Basic Information](#)
- [Enable/Disable Endpoints](#)
- [Add/Remove Alarms for Endpoints](#)
- [Delete Endpoints](#)

Modify Basic Information

If you only need to modify the name and description of an endpoint, you can do so on the **Endpoint** page by clicking **Actions > Edit Name and Description**.

If you need to modify the configuration information of an endpoint, you can do so on the **Endpoint** page by clicking **Actions > Modify Configuration**.

Enable/Disable Endpoints

To enable or disable one or more endpoints to prevent unnecessary personnel from receiving alarm messages and ensure that relevant personnel can receive alarm information in a timely manner to take necessary measures, you can select the desired endpoints on the **Endpoint** page and then click **Enable** or **Disable**.

Add/Remove Alarms for Endpoints

To add or remove alarms for an endpoint to ensure it only receives relevant alarm information and avoid unnecessary disturbances, you can select an endpoint on the **Endpoint** page, then click **Actions > Add Alarm/Remove Alarm** and choose the alarms you wish to add or remove.

Delete Endpoints

To delete an existing endpoint, you can select the endpoint you wish to delete on the **Endpoint** page, then click **Actions > Delete** to remove it.

**Note:**

You cannot delete system endpoints.

12.2.3 Message Template

12.2.3.1 New Message Template

Message templates are text templates used by alarms to push notification messages to endpoints.

You can specify a default message template for each type of endpoint, and alarm messages will be sent using the format of the default template.

Prerequisites

Before creating message templates for DingTalk, Lark, WeCom, or Microsoft Teams, make sure NexaSphere has an installed Advanced Edition license, and that the license is in a valid state.

Procedure

1. In the navigation pane, choose **O&M Management > Message Template**.
2. On the **Message Template** page, click **New Message Template**.
3. In the **New Message Template** dialog, set the following parameters:

Basic Information

- **Name:** Enter the name for the message template.
- **Description:** Provide a description for the message template.

Template Information

When the type is selected as Email, DingTalk, Lark, WeCom, HTTP Application, or Microsoft Teams, configure the template as follows:

- **Type:** Select the message template type, choosing from Email, DingTalk, Lark, WeCom, HTTP Application, or Microsoft Teams.
- **Alarm Type:** Select resource alarms or event alarms.
- **Alarm Message Title:** The title displayed in the alarm message.
- **Alarm Message Text:** The text displayed in the alarm message.
- **Recovery Message Title:** The title of the recovery message sent when a monitored resource's status returns to normal. This parameter is only supported for resource alarms.
- **Recovery Message Text:** The text of the recovery message. This parameter is only supported for resource alarms.

- **Default Template:** Unchecked by default. Checking this option sets the current template as the default template.

When the type is selected as SMS, configure the template as follows:

- **Type:** Select SMS as the message template type.
- **Signature:** Enter the SMS signature name applied for on 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.
- **Default Template:** Set this template as the default template. After setting, all SMS messages will be sent in this template format.



Note:

1. To create an Email or Lark message template, follow the Text syntax.
2. To create a DingTalk or WeCom message template, follow the Markdown syntax
3. To create a HTTP Application message template, follow the JSON syntax
4. To create a Microsoft Teams message template, follow the Webhook syntax requirements listed on the Microsoft Teams official website
5. To create a SMS message template, you need to apply for third-party SMS signatures templates in advance. Currently, you can use Alibaba Cloud SMS service. Any template changes require re-applying through the third-party service.

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

12.2.3.2 Manage Message Template

- [Modify Basic Information](#)
- [Set Default Message Template / Unset Default](#)
- [Delete Message Templates](#)

Modify Basic Information

If you only need to modify the name and description of a message template, you can do so on the **Message Template** page by clicking **Actions > Edit Name and Description**.

If you need to modify other configurations of the message template, including basic information and template information, you can do so on the **Message Template** page by clicking **Actions > Modify Configuration**.

Set Default Message Template / Unset Default

If you have added multiple message templates and need to specify one as the default message template, alarm messages will be sent using the format of the default template. You can set a template as the default on the **Message Template** page by clicking **Actions > Set as Default**.

If you need to unset the default message template, you can do so on the **Message Template** page by clicking **Actions > Unset Default**.

Delete Message Templates

To delete one or more message templates, you can do so on the **Message Template** page by clicking **Bulk Actions > Delete**.

12.2.4 Alarm

12.2.4.1 Alarm Rules

Before creating a new alarm, you may want to familiarize yourself with the alarm rules. This section will demonstrate how to configure alarm rules for resource-based and event-based alarms.

Resource Alarm Rules

Parameter		Description	Example
Resource Type		The type of resource monitored by the alarm.	Virtual Machine
Metric Item		Types and names of various monitoring metrics	CPU Usage
Resource		The specific resource object monitored by the alarm.	/
Alarm Trigger Rule	Comparison Relation	Defines how the detected metric value compares to the threshold. Comparison relations include >, ≥, <, and ≤.	>
	Threshold	The threshold value and unit that triggers an alarm.	85%
	Duration	The duration for which the alarm must continuously trigger before sending an alarm message.	5 minutes

Parameter		Description	Example
		Durations include 30 seconds, 1 minute, 5 minutes, 10 minutes, 30 minutes, 1 hour, custom.	
Alarm Interval		When an alarm is triggered, it notifies at a specific interval. Intervals include once only, every 1 hour, custom.	Every 1 hour
Emergency Level		Alarm message levels include Emergent, Major , Info.	Major
Recovery Notification		Sends a recovery notification when the monitored resource returns from an alarm state to a normal state.	/

Event Alarm Rules

Parameter		Description	Example
Resource Type		The type of resource monitored by the alarm.	Data Storage
Metric Item		Name of the monitored event or metric.	Data Storage Disconnected
Emergency Level		Alarm message levels include Emergent, Major , Info.	Emergent

12.2.4.2 New Resource Alarm

Resource alarms are used to monitor time-series data of resources within the platform. For example, you can set an alarm for the CPU usage of a virtual machine. If the CPU usage exceeds 80% and persists for 5 minutes, it will trigger a system alarm.

Prerequisites

- NexaSphere provides system parameter functionality to globally control the default behavior of platform settings. You can customize parameters related to alarms within the system parameters. For more information, see [System Parameters](#).
- Some alarm items require VMTools for monitoring and alerting. For more information about VMTools, see [Virtual Machine VMTools](#).

Procedure

- In the navigation pane, choose **O&M Management > Alarm > Resource Alarm**.
- On the **Resource Alarm** page, click **New Resource Alarm**.

3. In the **New Resource Alarm** dialog, set the following parameters:

Basic Information:

- **Name:** Enter the name for the resource alarm.
- **Description:** Provide a description for the resource alarm.

Configuration Information:

- **Resource Type:** Select the type of resource.
- **Metric Item:** Choose alarm items based on the selected resource type. Some alarm items require selecting a resource and setting corresponding alarm trigger rules.
- **Alarm Interval:** Choose the alarm interval type, including one-time and repeated alarms.
- **Emergency Level:** Different levels of alarms will issue messages according to their severity, including Emergent, Major, and Info.
- **Alarm Recovery Notification:** Disabled by default. If enabled, you will receive a recovery notification when any monitored resource returns to a normal state from an alarm state.

Recovery notifications are sent using the default message template, but you can also customize the message template. For more information, see [New Message Template](#).

- **Endpoint:** After the alarm is triggered, the alarm message will be pushed to the specified endpoint.

The system provides the default endpoint. You can also create custom endpoints. For more information, see [New Endpoint](#).

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

12.2.4.3 New Event Alarm

Event alarms are used to monitor predefined events within the platform. For example, a host disconnection event alarm will trigger a system alarm when the host becomes unreachable.

Prerequisites

- NexaSphere provides system parameter functionality to globally control the default behavior of platform settings. You can customize parameters related to alarms within the system parameters. For more information, see [System Parameters](#).
- Some alarm items require VMTools for monitoring and alerting. For more information about VMTools, see [Virtual Machine VMTools](#).

Procedure

1. In the navigation pane, choose **O&M Management > Alarm > Event Alarm**.
2. On the **Event Alarm** page, click **New Event Alarm**.
3. In the **New Event Alarm** dialog, set the following parameters:
 - **Resource Type**: Select the type of resource.
 - **Metric Item**: Choose alarm items based on the selected resource type.
 - **Emergency Level**: Different levels of alarms will issue messages according to their severity, including Emergent, Major, and Info.
 - **Endpoint**: After the alarm is triggered, the alarm message will be pushed to the specified endpoint.

The system provides the default endpoint. You can also create custom notification objects.

For more information, see [New Endpoint](#).

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

12.2.4.4 Manage Alarm

- [Modify Basic Information](#)
- [Enable/Disable Alarms](#)
- [Add/Remove Endpoints for Alarms](#)
- [Delete Alarms](#)

Modify Basic Information

If you only need to modify the name and description of an alarm, you can do so on the **Alarm** page by clicking **Actions > Edit Name and Description**.

If you need to modify the basic information and configuration of a resource alarm, or the configuration information of an event alarm, you can do so on the **Alarm** page. Select the target resource alarm or event alarm, then click **Actions > Modify Configuration** to make changes.

Enable/Disable Alarms

To enable or disable one or more alarms, you can do so on the **Alarm** page by selecting the alarms and then clicking **Enable** or **Disable**.

Add/Remove Endpoints for Alarms

To add or remove endpoints for an alarm, ensuring that endpoints only receive relevant alarm information and avoid unnecessary disturbances, you can do so on the **Alarm** page by clicking **Actions > Add Endpoint/Remove Endpoint**. Select the endpoints you wish to add or remove.

Delete Alarms

To delete one or more alarms, you can do so on the **Alarm** page by selecting the alarms and then clicking **Delete**.

**Note:**

You cannot delete default alarms.

12.2.5 Alarm Message

12.2.5.1 View Alarm Messages

Context

Triggered alarms are visible in several locations throughout the platform.

- View from the alarm message page: The alarm messages page presents the overall platform alarm information on a single dashboard. You can compare and view alarm messages based on different dimensions, helping you intuitively and comprehensively understand the platform's resource status and identify potential issues and bottlenecks.
- View from the resource's monitoring tab: You can focus on a specific resource to view its alarm messages, gaining a more detailed understanding of the alarm conditions for that resource. This allows for more targeted optimization and adjustments.
- View from the bottom task and alarm pane: You can check alarm messages from the pane at the bottom of the platform. The pane displays up to 50 recent alarm messages, and you can also click **More** to navigate to the alarm messages page.

Procedure

1. In the navigation pane, choose **O&M Management > Alarm Message**.
2. View the triggered alarm messages.

The alarm message page consists of weekly alarm statistics, weekly alarm distribution, and an alarm message list.

- **Alarm Statistics in Recent 1 Week:** Displays alarm statistics for the past 7 days in a bar chart format, with a sampling interval of 8 hours. Hover over the bar chart to view the number of alarms at different levels.
- **Alarm Distribution in Recent 1 Week:** Shows the percentage of resource alarms within the last 7 days using a bar chart. Hover over the bar chart to view the number of alarms for different types of resources.

- **Alarm Message List:** Displays up to 1,000 alarm messages in a list format. You can filter the display by resource type and time.

12.2.5.2 Acknowledge Alarm Messages

Acknowledging an alarm lets other users know that you are taking ownership of the issue.

Administrators acknowledge alarm messages to make it easier for O&M personnel to identify and respond to them promptly, ensuring no critical alarm information is missed.

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**.
 1. Acknowledged alarms will not be displayed in the triggered alarms tab but can be viewed in the all alarms tab.
 2. After acknowledging, if the alarm issue is not resolved promptly, the alarm system will continue to trigger and push messages according to the rules. To avoid repeated notifications, you can set a silence period as needed.

12.2.5.3 Set Silence Period for Alarm Messages

If you need to temporarily suspend the push of a specific alarm message for a certain period, you can set a silence period for it. During the silence period, the alarm message will not be pushed.

Once the silence period ends, if the alarm remains triggered, the alarm message will be pushed again.

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.

12.2.5.4 Restore Alarms

To resume the push of alarm messages during the silence period, you can manually restore the alarm.

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 **Restore Alarm**.

12.3 Task

Operation Task

Operation Task: An operation task is a chronological record of operations on the specified objects and their operation results.

The operation task displays the historic and ongoing operations in the platform.

- You can select a time period to view operation tasks within that range. Options include: Last 3 days, Last 7 days, Last 1 month, and Custom.
- You can search for operation tasks by entering the task description or login IP.
- You can filter operation tasks by the results. Options include: Succeeded, Failed, Canceled, Canceling, Abnormal, Ongoing, Timeout, Suspended, and Unknown.
- You can sort operation tasks by start time, completion time, or consumed time.
- You can export operation tasks in CSV format.
- You can adjust the number of operation tasks displayed per page. Options include: 10, 20, 50, 100 items/page. The interface supports pagination.
- You can cancel, pause, or resume ongoing operation tasks.
- The platform retains operation tasks for 90 days by default. You can customize the maximum retention period for operation tasks. For more information, see [System Parameters](#).

HA Task

HA Task: HA task logs generated when the platform executes high availability procedures in accordance with the enabled HA policy.

The HA task displays all virtual machine HA logs in the platform.

- You can select a time period to view HA tasks within that range. Options include: Last 7 days, Last 1 month, and Custom.
- You can search for HA tasks by entering the VM name, pre host, or destination host.
- You can filter HA tasks by the results. Options include: Succeeded and Failed.
- You can sort HA tasks by start time or completion time.
- You can export HA tasks in CSV format.
- You can adjust the number of HA tasks displayed per page. Options include: 10, 20, 50, 100 items/page. The interface supports pagination.

Scheduling Task

Scheduling Task: Scheduling task logs generated when the platform executes dynamic resource scheduling operations after the cluster DRS is enabled.

The scheduling task displays the virtual machine automatic scheduling logs triggered by the management node.

- You can select a time period to view scheduling tasks within that range. Options include: Last 3 days, Last 7 days, Last 1 month, and Custom.
- You can search for scheduling tasks by entering the migration object or UUID.
- You can adjust the number of scheduling tasks displayed per page. Options include: 10, 20, 50, 100 items/page. The interface supports pagination.
- The platform retains scheduling tasks for 90 days by default. You can customize the maximum retention period for scheduling tasks. For more information, see [System Parameters](#).

12.4 Event

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. Event allows for continuous monitoring and retention of all platform activities, including the use and access of the platform via the console, API services, and developer tools.

Resource Action Events

Resource action events provide information about API calls made during operations on resources within the platform.

- Click the API name to view event details, including basic information, API requests, API response, and more.
- You can select a time period to view events within that range. Options include: Last 3 days, Last 7 days, Last 30 days, and Custom.
- You can search for resource action events by entering the API name, resource UUID, name, resource type, or operator IP.
- You can filter events by the results. Options include: Succeeded and Failed.
- You can export events in CSV format.
- You can adjust the number of events displayed per page. Options include: 10, 20, 50, 100 items/page. The interface supports pagination.

- The platform retains all event records. You can customize the maximum retention period for events. For more information, see [System Parameters](#).

Login Operation Events

Login operation events provide event information for login and logout API operations.

- Click the API name to view event details, including basic information, API requests, API response, and more.
- You can select a time period to view events within that range. Options include: Last 3 days, Last 7 days, Last 30 days, and Custom.
- You can search for login operation events by entering the API name, login IP, or browser.
- You can filter events by the results. Options include: Succeeded and Failed.
- You can export events in CSV format.
- You can adjust the number of events displayed per page. Options include: 10, 20, 50, 100 items/page. The interface supports pagination.
- The platform retains all event records. You can customize the maximum retention period for events. For more information, see [System Parameters](#).

12.5 Log Collection

12.5.1 Collect Logs

Prerequisites

- Ensure sufficient storage space. Insufficient storage space will cause collection failures.
- The platform supports a maximum of three collected logs. To collect new logs, delete existing ones as needed.
- To download collected logs via URL, make sure the download environment has network access to 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. In the **Collect Log** dialog, set the following parameters:
 - **Log Type**: Select the type of logs to collect.

- **All logs:** Collect all platform logs types with one click, including operation logs, event logs, management node logs, compute node logs, data storage logs, image storage logs, and database logs.
 - **Specific logs:** Collect only selected log types. For example, selecting compute node logs will collect logs from all compute nodes within the specified time range.
 - **Time Range:** Specify a time range to collect logs. Maximum duration: 5 days.
 - **Auto-Download:** Choose whether to automatically download the collected logs.
4. Review the configuration and click **OK**.

12.5.2 Manage Collected Logs

Procedure

1. In the navigation pane, choose **O&M Management > Log Collection**.
2. On the **Log Collection** page:
 - a) Click **Download** to download the collected logs.
 - b) Click **Delete** to delete the collected logs.
 - c) Click **Delete All Logs** to delete all collected logs with one click.

12.6 Tag Management

12.6.1 Overview

Tags serve as markers for resources. You can assign different tag definitions based on various business needs and purposes. By customizing and attaching tags to resources, you can quickly filter out the required resources and improve search efficiency.

User Permissions

Tags are divided into admin tags and user tags, and you can determine their type by the **owner**.

- **Admin Tags:** Created by the admin and owned by the admin.
- **User Tags:** Created by regular users and owned by regular users.



Note:

Changing the owner of a tag is not supported.

Admin can detach or delete tags created by regular users. When the ownership of a resource changes, all user tags on it will automatically be detached, while admin tags remain unaffected.

Resources That Support Tagging

Resource	Supports Tagging	Supports Tagging at Creation
Host	Yes	Yes
Virtual Machine	Yes	Yes
Bare Metal Instance	Yes	No

Tag Limitations

- Tags must be unique. Users with the same role can create only one tag with the same name and color.
- Tag names must not exceed 20 characters in length.
- You can attach admin tags to hosts, virtual machines, and bare metal instances.
- You can attach user tags to virtual machines.

12.6.2 Create a Tag

Create simple, minimalistic tags with different colors.

Procedure

1. In the navigation pane, choose **O&M Management > Tag**.
2. On the **Tag** page, click **New Tag**.
3. In the **New Tag** dialog, set the following parameters:
 - **Tag Preview**: Displays how your 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.
 - Tag names must not exceed 20 characters in length.
 - **Description**: Optional. Enter tag details.
 - **Color**: Select a tag color.
4. Review the configuration and click **OK**.

12.6.3 Attach/Detach a Tag

After you create tags, you can attach them to resources or detach them from resources.

Procedure

1. (Optional) Attach/detach tags from the resource perspective.
 - a) In the navigation pane, choose **Inventory > VM and Host/Bare Metal Management**.
 - b) Select a target host/virtual machine/bare metal instance from the resource tree.
 - c) On the resource details page, click **Actions > Tag and Attribute > Tag Management**.
 - d) In the **Tag Management** dialog, add or remove tags as needed.
 - e) Click **OK**.
2. (Optional) Attach/detach tags from the tag perspective.
 - a) In the navigation pane, choose **O&M Management > Tag**.
 - b) On the **Tag** page, click the name of a tag to enter its details page.
 - c) Click the **Associated Resource** tab.
 - d) On the **Associated Resource** tab, select resources to attach/detach, and then click **Attach/Detach**.

12.6.4 Delete a Tag

Procedure

1. In the navigation pane, choose **O&M Management > Tag**.
2. On the **Tag** page, select tags to delete.
3. Click **Actions > Delete**.
4. To confirm the operation, click **OK**.

12.6.5 Search Resources Using Tags

This section introduces how to quickly filter target virtual machines using tags.

Procedure

1. In the navigation pane, choose **Inventory > VM and Host**.
2. Select a host from the resource tree.
3. On the details page of the selected host, click **Virtual Machine** tab.
4. On the **Virtual Machine** tab, click **Tag**.
5. From the drop-down menu, select tags.

- Tags are categorized as admin tags and user tags.
- You can search for specific tags by entering the tag name.

6. Click **OK**.

12.7 Custom Attribute

12.7.1 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, set the following parameters:
 - **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 a attribute key name.
 - **Description**: Optional. Enter a description for the custom attribute.
 - **Attribute Value**: Enter a attribute value name.
4. Click **OK**.

12.7.2 Add and Edit Custom Attributes

After you create a custom attribute, the attribute key will be automatically associated with all resources of the selected type. You can then add attribute values to the each resource.

Context

This section introduces how to configure custom attributes for virtual machines.

Procedure

1. In the navigation pane, choose **Inventory > VM and Host**.
2. From the resource tree, select a target virtual machine.
3. On the target virtual machine's details page, click **Actions > Tag and Attribute > Set Custom Attribute**.
4. In the **Set Custom Attribute** dialog, configure attribute value as needed.
 - If you want to add an attribute value to an existing attribute key, select or enter the attribute value in the **Attribute Value** column.

- If you want to create a new custom attribute, click **Add** and enter an attribute key and attribute value.
- If you batch configure custom attributes for multiple virtual machines, batch configuration will overwrite the corresponding attribute values for all selected resources. If left blank, each resource's existing attribute values will be preserved.

5. Click **OK**.

12.7.3 Manage a Custom Attribute

Procedure

1. In the navigation pane, choose **O&M Management > Custom Attribute**.
2. On the **Custom Attribute** page, select a target custom attribute.
 - Click **Actions > Modify Configuration** to edit the attribute key and description of a custom attribute.
 - Click **Actions > Add Attribute Value** to add a new attribute value to the custom attribute.
 - Click **Actions > Delete** to delete a custom attribute.
 - Click the name of a custom attribute to enter its details page. You can have a fine-grained management of attribute values, including adding, deleting, and viewing associated resources.

12.7.4 Search Resources Using Custom Attributes

Context

This introduces how to quickly filter target virtual machines using custom attributes.

Procedure

1. In the navigation pane, choose **Inventory > VM and Host**.
2. From the resource tree, select a host.
3. On the details page of the selected host, click **Virtual Machine** tab.
4. On the **Virtual Machine** tab, click **Custom Attribute**.
5. From the drop-down menu of custom attribute, select attribute keys and attribute values.

When the attribute value is empty, the platform show resources with unset attribute values under this attribute key.

6. Click **OK**.

12.8 Migration Service

The migration service enables complete migration of virtual machine systems and data from other virtualization platforms to the current platform. Supported source platforms include VMware vSphere and other KVM-based cloud platforms. A step-by-step wizard guides you through the end-to-end migration process.

12.8.1 Deploy Migration Service

12.8.1.1 Step 1: Upload Migration Service Package

After uploading the installation package, the system automatically decompresses it and extracts the required files.

Prerequisites

- Ensure that the migration service license is installed on the platform and that you have sufficient license capacity.
- Prepare the migration service installation package in advance.
- Ensure that an image storage is added to the platform and that the image storage has sufficient storage space.

Procedure

1. In the navigation pane, select **O&M Management > Migration Service**.
2. On the **Migration Service** page, click **Upload**.
3. In the **Upload Migration Service Package** dialog, set the following parameters:
 - **Upload By**: Select the method for uploading the installation package, including URL and local file.
 - **URL/Installation Package**: Enter a URL or upload the installation package file.
 - **Storage Path**: Specify a directory path to store the uploaded installation package and its decompressed files.



Note:

- Ensure that the directory has sufficient available space. Otherwise, the installation fails.
- Do not use system directories such as `/`, `/dev`, `/proc`, `/sys`, `/usr/bin`, `/bin`, or `/opt`. Using these directories might cause server issues.

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

12.8.1.2 Step 2: Deploy Migration Service

Select the appropriate specification based on your business requirements to deploy the migration service. Once deployed, you can use the migration service to migrate virtual machines to the target platform.

Procedure

1. On the **Migration Service** page, click **Install**.
2. In the **Deploy Migration Service** dialog, set the following parameters:
 - **Migration Service Version**: Displays the version of the service to be installed and deployed.
 - **Data Center**: The migration service is deployed through the virtual machine data gateway. Select a data center.
 - **Run Location**: Select where the service runs.
 - **CPU**: Default is 8 cores. Set the number of CPU cores.
 - **Memory**: Default is 8 GB. Set the memory size.
 - **Storage Location**: Specify the data storage.
 - **Storage Pool**: When the storage location is distributed storage, you must specify the storage pool.
 - **Hard Disk 1**: Default is 100 GB. Set the system disk size.
 - **Migration Network**: Specify a distributed port group to use as the migration network. Ensure that the selected distributed port group has network connectivity with both the source platform and the destination platform to guarantee data transfer.
 - **IPv4 Address**: Specify the IP address.



Note:

The IP address allocation method depends on the configuration of the selected migration network:

- When DHCP is enabled on the migration network, the system automatically assigns an IP address.
- When DHCP is disabled on the migration network, you need to specify an IP address.
- When IPAM is enabled on the migration network, manually specified IP address must be within the selected network segment.

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

12.8.2 Manage Migration Service

12.8.2.1 Upgrade the Migration Service

Prerequisites

- Ensure that no migration tasks are in progress. Otherwise, you cannot upgrade the migration service.
- Ensure that the image storage has sufficient storage space before you upload the migration service upgrade package.

Procedure

1. In the navigation pane, select **O&M Management > Migration Service > Migration Overview**.
2. On the **Migration Overview** page, click **Upgrade**.
3. In the **Upgrade Migration Service** dialog, set the following parameters:
 - **Migration Service Version**: Displays the currently deployed service version.
 - **Upload Method**: Select the method for uploading the installation package, including URL and local upload.
 - **URL/Installation Package**: Enter a URL or upload the installation package file.
 - **Storage Path**: Specify a directory path to store the uploaded upgrade package and its decompressed files.



Note:

- Ensure that the directory has sufficient available space. Otherwise, the installation fails.
- Do not use system directories such as `/`, `/dev`, `/proc`, `/sys`, `/usr/bin`, `/bin`, or `/opt`. Using these directories might cause server issues.

4. Click **OK**.

What's next

If the upgrade fails, you can upgrade again. You can use the existing upgrade package or upload a new one.

12.8.2.2 Enable/Disable the Migration Service

Prerequisites

Ensure that no migration tasks are in progress. Otherwise, you cannot disable the migration service.

Procedure

1. In the navigation pane, select **O&M Management > Migration Service > Service Management**.
2. On the **Service Management** page, select the target migration service, and then click **Actions > Enable/Disable**.

After you disable the migration service, all migration service-related features will become unavailable. Proceed with caution.

12.8.2.3 Delete the Migration Service

Prerequisites

- Ensure that no migration tasks are in progress. Otherwise, you cannot delete the migration service.
- You can delete the migration service only when it is disabled.

Procedure

1. In the navigation pane, select **O&M Management > Migration Service > Service Management**.
2. On the **Service Management** page, select the target migration service, and then click **Actions > Delete**.



Note:

After deletion, the system will permanently remove all migration service files and configurations, and immediately release the CPU, memory, and IP address resources. Proceed with caution.

3. After you read and confirm the risk information, click **OK**.

12.8.3 Manage Migration Resources

12.8.3.1 Add a Source Platform

Procedure

1. In the navigation pane, select **O&M Management > Migration Resource > Source Platform**.
2. On the **Source Platform** page, click **Add Source Platform**.
3. In the **Add Source Platform** dialog, set the following parameters:
 - **Name**: Set a name for the source platform.

- **Data Gateway Configuration Method:** Select the configuration method, including selecting an existing data gateway or entering a new data gateway IP address.
- **Data Gateway:** Select the data gateway.
- **Platform Type:** VMware.
- **vCenter/ESXi Address:** Enter the address of the VMware environment.
- **Username:** Enter the login username.
- **Password:** Enter the password for the username.
- **Connectivity Test:** Click **Test** to verify network connectivity.

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

12.8.3.2 Delete a Source Platform

Prerequisites

Ensure that no migration tasks are in progress. Otherwise, you cannot delete the source platform.

Procedure

1. In the navigation pane, select **O&M Management > Migration Resource > Source Platform**.
2. On the **Source Platform** page, select the target source platform, and then click **Actions > Delete**.
3. Click **OK**.

12.8.3.3 Add a Data Gateway

Procedure

1. In the navigation pane, select **O&M Management > Migration Resource > Data Gateway**.
2. On the **Data Gateway** page, click **Add Data Gateway**.
3. In the **Add Data Gateway** dialog, set the following parameters:
 - **Name:** Set a name for the data gateway.
 - **Addition Method:** Select the method for adding the data gateway. You can create a new virtual machine or use an existing IP address.

If you select **New Virtual Machine**:

- **Data Center:** Select the data center where the virtual machine resides.
- **Location:** Select where the virtual machine runs.
- **CPU:** Default is 8 cores. Set the number of CPU cores.
- **Memory:** Default is 8 GB. Set the memory size.

- **Storage Location:** Specify the data storage.
- **Disk 1:** Default is 100 GB. Set the system disk size.
- **Network:** Specify a distributed port group. Ensure that the selected distributed port group has network connectivity with both the source platform and the destination platform to guarantee data transfer.
- **IPv4 Address:** Specify the IP address.

**Note:**

The IP address allocation method depends on the configuration of the selected migration network:

- When DHCP is enabled on the migration network, the system automatically assigns an IP address.
- When DHCP is not enabled on the migration network, you must manually specify an IP address.
- When IP address management is enabled on the migration network, if you manually specify an IP address, ensure that the address is within the selected network segment.

If you select **Existing IP**:

- **IP Address:** Enter the IP address.
- **Host Name:** After the IP address is connected, the corresponding name is automatically retrieved.
- **Connectivity Check:** Click **Check** to verify network connectivity.

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

12.8.3.4 Delete a Data Gateway

Prerequisites

Ensure that no ongoing migration tasks are associated with the data gateway. Otherwise, you cannot delete the data gateway.

Procedure

1. In the navigation pane, select **O&M Management > Migration Resource > Data Gateway**.
2. On the **Data Gateway** page, select the target data gateway, and then click **Actions > Delete**.
3. After you read and confirm the risk information, click **OK**.

12.8.4 Manage Migration Tasks

12.8.4.1 Create a Migration Task

Prerequisites

Ensure that you have sufficient license capacity for the migration service.

Procedure

1. In the navigation pane, select **O&M Management > Migration Task > All Tasks**.
2. On the **All Tasks** page, click **Create Migration Task**.
3. In the **Create Migration Task** dialog, complete the source platform configurations:
 - **Source Platform:** Select the source platform.
 - **Source VM:** Add the source virtual machine you want to migrate.



Note:

The disk capacity of the source VM cannot exceed the limit defined by the license type. Additional capacity requires a separate license.

- Trial License: Disk capacity limit is 1 TB.
 - Production License: Disk capacity limit is 5 TB.
- **Migration Authorization Type:** Select the license type.
 4. Click **Next**.
 5. Configure the target platform.
 - a) Complete the virtual machine configuration:
 - **Data Center:** Select the data center where the virtual machine resides after migration.
 - **Cluster:** Select the cluster where the virtual machine resides after migration.
 - **Host:** Select the host where the virtual machine resides after migration.
 - **Storage Location:** Select the data storage that the virtual machine uses after migration.
 - **Bus Type:** Select the bus type for the virtual machine hard disk device.
 - **CPU Mode:** Select the CPU mode for the virtual machine after migration.
 - **Destination Network Port Group:** Select the distributed port group that the virtual machine uses.
 - **CIDR:** Select the CIDR.
 - **Network Adapter Configuration:** Select whether to clone the source virtual machine's network configuration and MAC address.

- **Clone Network Configuration:** After cloning, the target virtual machine will retain the same IP address as the source VM by default. The netmask and gateway must be manually entered.
- **Clone MAC Address:** The target MAC address of the virtual machine remains the same as that of the source virtual machine.

**Note:**

- When you select Clone Network Configuration, the platform assigns a temporary IP address to the virtual machine based on network properties to ensure data synchronization. You can also manually specify an IP address on the virtual machine configuration confirmation page.
 - When you select Clone Network Configuration, power off the source virtual machine before the formal cutover to avoid IP address conflicts.
- **Destination IP Address:** Specify the destination IP address for the virtual machine.
 - **Temporary Network Port Group:** When you select **Clone Network Configuration**, you must configure a temporary network port group. A DHCP-enabled port group is recommended. By default, the target network port group is used. If the target port group does not have DHCP enabled, you must manually assign a temporary IP address.
 - **Security Group:** Select a security group.
- b) Complete the cutover configuration by referring to the following example:
- **Immediate Cutover:** When enabled, the system automatically performs the cutover task immediately after the initial data synchronization is completed.
 - **Create Rollback Snapshot:** When enabled, the system creates a rollback snapshot for the target VM when the cutover task is initiated.

Purpose of the rollback snapshot:

- When enabled: If you select "Delete cutover program only" when deleting the migration task, the system automatically rolls back the target VM or disk to the snapshot state.
- When disabled: The "Delete cutover program only" option will not be available when deleting the migration task.
- **Auto Shut Down Source VM:** When enabled, the system shuts down the source VM before executing the cutover task.

- **Auto Install VMTools:** When enabled, the system automatically installs VMTools after the target VM operating system conversion is complete.

6. Click **Next**.

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

12.8.4.2 Synchronize a Migration Task

Prerequisites

Ensure that the migration task status is "Synchronization Stopped", "Synchronization Failed", or "Ready for Cutover". Otherwise, you cannot perform the synchronization operation.

Procedure

1. In the navigation pane, select **O&M Management > Migration Task > All Tasks**.
2. On the **All Tasks** page, select the migration task you want to synchronize, and then click **Actions > Synchronize**.
3. In the **Synchronize Migration Task** dialog, select the synchronization mode.

The following two synchronization modes are supported:

- **Incremental Sync:** Only synchronizes data that has changed since the last synchronization.
- **Full Sync:** Synchronizes all data of the virtual machine. For migration tasks that have completed initial synchronization, the system performs a comparison scan before executing a full sync. This process may take a significant amount of time.

4. Click **OK**.

12.8.4.3 Create a Test Virtual Machine

Prerequisites

Ensure that the migration task status is not "Cutover in Progress". Otherwise, you cannot create a test virtual machine.

Procedure

1. In the navigation pane, select **O&M Management > Migration Task > All Tasks**.
2. On the **All Tasks** page, select the target migration task, and then click **Actions > Test > Create Test Virtual Machine**.
3. In the **Create Test Virtual Machine** dialog, set the following parameters:
 - **CPU:** Set the number of CPU cores for the virtual machine.
 - **Memory:** Set the memory size for the virtual machine.

- **Network:** Select the distributed port group that the virtual machine uses.
- **CIDR:** The system automatically retrieves the CIDR based on the selected network.
- **IPv4 Address:** The IP address allocation method depends on the configuration of the selected network.
 - If DHCP is enabled on the selected distributed port group, an IPv4 address is automatically assigned by default. You can also manually specify an IP address, and the specified IP address must be within the CIDR.
 - If DHCP is not enabled on the selected distributed port group, you must manually specify an IP address, and the specified IP address must be within the CIDR.
 - If IPAM is not enabled on the selected distributed port group, the CIDR parameter is not displayed. You must manually specify an IP address.
- **System Conversion Policy:** Specify the system conversion policy before the virtual machine starts.
 - **Perform System Conversion:** Performs a full system conversion, including: driver replacement/injection, network adapter name and IP address configuration for the target VM, and post-script execution.
 - **Perform System Conversion Without Loading Target Drivers:** Skips driver replacement. Only configures the target VM's network adapter name and IP address, and executes post-scripts.
 - **Skip System Conversion:** No driver replacement, no network configuration (adapter name/IP address), and no post-script execution. You will need to log in to the target VM and manually perform these configurations.
- **Data Recovery Point:** Use a snapshot or copy from a specific point in time to perform virtual machine recovery testing.

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

12.8.4.4 Perform Cutover

Prerequisites

Ensure that the migration task status is "Ready for Cutover". Otherwise, you cannot perform the cutover operation.

Procedure

1. In the navigation pane, select **O&M Management > Migration Task > All Tasks**.
2. On the **All Tasks** page, select the target migration task, and then click **Actions > Cutover**.

After cutover starts, virtual machine services are formally switched to the target platform. If the migration task has "Auto Shut Down Source VM" enabled, the system powers off the source virtual machine according to the shutdown method you selected.

3. After you read and confirm the risk information, click **OK**.

12.8.4.5 Delete a Migration Task

Prerequisites

- When you delete multiple migration tasks, you can directly delete migration tasks whose status is "Synchronizing", "Synchronization Stopped", "Synchronization Failed", or "Ready for Cutover".
- When you delete multiple migration tasks, for migration tasks whose status is "Cutover in Progress", "Cutover Successful", or "Cutover Failed", you can select a deletion policy as needed.

Procedure

1. In the navigation pane, select **O&M Management > Migration Task > All Tasks**.
2. On the **All Tasks** page, select the target migration task, and then click **Actions > Delete**.
3. In the **Deletion Policy** dialog, select a deletion policy.
 - **Retain target Instance and Delete Migration Task**: Only deletes the migration task record. The target instance is retained. You can reconfigure the source virtual machine and perform a new migration task later. This operation does not cause data or state loss, but you must manage the target instance yourself.
 - **Delete target Instance and Delete Migration Task**: Deletes the migration task record and its associated target instance. You can reconfigure the source virtual machine and perform a new migration task later.
 - **Delete Cutover Program Only**: Removes only the cutover program. This policy is available only when "Create Rollback Snapshot" is enabled. The system automatically rolls back the target instance to the state before cutover, and all changes made after cutover are lost. This operation is irreversible. Use with caution.
4. Click **OK**.
5. After you read and confirm the risk warning, click **OK**.

13 System Management

13.1 Identity and Access Management

13.1.1 Overview

The Identity and Access Management module of NexaSphere provides unified user identity management and access control. It supports centralized management of regular users, configuration of a unified authentication system for single sign-on with NexaSphere, and management of access permissions for all users to platform resources.

Centralized User Management

It supports the unified creation and management of users and user groups.

Precise Role-Based Access Control

Different role permissions can be granted to different users and user groups, allowing you to precisely control the operations that specific users or user groups can perform on particular resources, thereby assisting in maintaining the security of the environment.

Integration with Unified Authentication Systems

It supports configuring a unified authentication system based on OIDC, AD, LDAP protocols for single sign-on (SSO). This allows direct use of users from the unified identity authentication system without the need to create additional users, enhancing management efficiency and reducing security risks.

13.1.2 Preparation

Before using the NexaSphere identity and access management features, ensure that the platform version and license authorization meet the requirements.

- To use the single sign-on system, roles, and user group features, ensure that the NexaSphere is installed with a valid Advanced Edition license.

13.1.3 Single Sign-On

NexaSphere provides unified identity authentication login services, supporting seamless access to the unified authentication login system. Corresponding unified authentication users can log in directly to the virtualization platform and conveniently 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, allowing clients to verify user identities and obtain basic user configuration information. Through the OIDC authentication server, user information can be synchronized to the platform according to mapping rules, and OIDC authentication system 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 a standalone and standardized login authentication system for increasingly diverse enterprise office applications. Through the AD authentication server, AD users can be synchronized to the virtualization platform, supporting direct login to the platform using specified AD login attributes.
- **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 the LDAP authentication server, LDAP users can be synchronized to the virtualization platform, supporting direct login to the platform using specified LDAP login attributes.

13.1.3.1 Add OIDC SSO Server

Prerequisites

Make sure NexaSphere is installed with a valid Advanced Edition license.

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, set the following parameters:

Basic Information

- **Name:** Set a name for the unified authentication server.
- **Description:** Optionally fill in a description for the unified authentication server.
- **Type:** Select **OIDC**.

Configuration Information

- **Client ID:** The unique identifier assigned to the platform by the unified authentication system.
- **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 and description.
 - **Username:** Maps the virtualization platform user name to a specific attribute of users in the unified authentication server.
 - **Description:** Optional, maps the platform user description to a specific attribute of users in the unified authentication server.

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

13.1.3.2 Add AD SSO Server

Prerequisites

Make sure NexaSphere is installed with a valid Advanced Edition license.

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, set the following parameters:

Basic Information

- **Name:** Set a name for the unified authentication server.
- **Description:** Optionally fill in a description for the unified authentication server.
- **Type:** Select **AD**.

Server Information

- **SSL/TLS Encryption:** Choose whether to enable SSL/TLS encryption. This is enabled by default.

When enabled, port 636 is used by default, with support for custom modifications. When disabled, port 389 is used by default, with support for custom modifications.

- **Primary Server IP/Domain:** Enter the primary server IP address or domain along with the corresponding port.

- **Backup Server IP/Domain:** Enter the backup server IP address or domain along with the corresponding port.

Configuration Information

- **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.
- 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**.

13.1.3.3 Add LDAP SSO Server

Prerequisites

Make sure NexaSphere is installed with a valid Advanced Edition license.

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, set the following parameters:

Basic Information

- **Name:** Set a name for the unified authentication server.

- **Description:** Optionally fill in a description for the unified authentication server.
- **Type:** Select **LDAP**.

Server Information

- **SSL/TLS Encryption:** Choose whether to enable SSL/TLS encryption. This is enabled by default.

When enabled, port 636 is used by default, with support for custom modifications. When disabled, port 389 is used by default, with support for custom modifications.

- **Primary Server IP/Domain:** Enter the primary server IP address or domain along with the corresponding port.
- **Backup Server IP/Domain:** Enter the backup server IP address or domain along with the corresponding port.

Configuration Information

- **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**.

13.1.3.4 Manage SSO Server

Procedure

1. In the navigation pane, choose **System Management > Single Sign-On**.
2. On the **Single Sign-On** page, perform the following steps as required:
 - 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.

13.1.4 Role Management

A role is a collection of permissions that, when granted to users and user groups, enables them to invoke related APIs for resource operations. NexaSphere adopts a Role-Based Access Control (RBAC) authorization model, defining resource permissions based on the user's job function (role). Through roles, you can achieve fine-grained control over user permissions.

13.1.4.1 System Predefined Roles

NexaSphere provides predefined roles as shown in the following table.

User Type	Role Name	Description
Admin User	System Admin	Manages daily system operations and maintenance.
	Security Admin	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

User Type	Role Name	Description
Regular User	VM User	Supports regular users in creating virtual machines and basic VM management.

13.1.4.2 Create 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

- **Name:** Set a name for the role.
- **Description:** Fill in a description for the role as needed.

Permission Configuration

Select the interface permissions you want to grant to this role as required. There may be dependencies between different interface permissions. It is recommended to use platform predefined roles or select all interface permissions.

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

13.1.4.3 Clone 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**.

13.1.4.4 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.

13.1.4.5 Delete Role

Prerequisites

- Make sure NexaSphere is installed with a valid Advanced Edition license.
- The selected role is not a system default role.
- The selected role has been detached from its associated users or user groups.

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**.

13.1.5 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.
- Users can be a member of one or more user groups.
- Users can be assigned one or more roles. When a user is assigned multiple roles, they will have the combined permissions of those roles. Additionally, once a user joins a user group, they will inherit the roles associated with that group, in addition to any roles they are already assigned.

13.1.5.1 New 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, set the following parameters:

Basic Information

- **Username:** Set a username for the regular user, which serves as the unique identifier for logging into the platform.
- **Description:** Optionally fill in a description for the user.
- **Password:** Set a login password for the user.

The password setting requirements can be adjusted by modifying the **Platform Login Password Strength**. For more information, see [Security Settings](#).

- **Confirm Password:** Re-enter the login password for confirmation.
- **Role:** Assign roles to the user. After binding, the user will have the permissions associated with the role.
- **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

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**.

13.1.5.2 Disable/Enable 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**.

13.1.5.3 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.

13.1.5.4 Change 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**.

13.1.5.5 Change a Regular 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, associated user groups, and shared resources will be automatically removed from the regular user.

13.1.5.6 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.

13.1.6 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 associate it with roles that involve storage resource permissions, then all users within this group will automatically inherit the role permissions from

UserGroup-1. If there's a new user who needs storage resource permissions, you can achieve the necessary permission allocation by adding this user to the UserGroup-1 user group. In case of changes in users, such as replacing old users with new ones, you don't need to modify permissions for each old user individually; instead, you can simply remove the old users from the user group.

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 only contain users, not other user groups.
- A user group can be assigned multiple roles. When a user group is assigned multiple roles, users within the group will inherit the combined permissions of those roles.

13.1.6.1 New 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

- **Name:** Set a name for the user group.
- **Description:** Optionally fill in a 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 assigning, all users within the group will inherit the permissions associated with these roles.

Share Resource

Share resources with the user group. After sharing, all users within the user group will have read access to the shared resources.

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

13.1.6.2 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.

13.1.6.3 Delete 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.

13.2 Security Access Settings

NexaSphere helps you improve access control security through IP blocklists and allowlists, certificate management, and security settings.

This chapter mainly covers the following topics:

- [IP Blocklist and Allowlist Management](#)

- [Certificate Management](#)
- [Security Settings](#)

13.2.1 IP Blocklist and Allowlist Management

IP blocklist and allowlist: By identifying and filtering visitor IPs, it intercepts access from specific IPs or allows access from specific IPs, further enhancing the access control security of NexaSphere.

- [Basic Information](#)
- [Add IP Blocklist or Allowlist](#)
- [Manage IP Blocklist and Allowlist](#)

Basic Information

- If no IP blocklists or allowlists have been added, requests from all IP addresses are allowed by default.
- If only an IP blocklist is added, IPs in the blocklist are denied access to the platform, while other IPs are allowed.
- If the same IP is added to both lists, the allowlist takes precedence over the blocklist, allowing requests from that IP.
- You cannot use the IP allowlist alone. Make sure to add at least one IP blocklist before you use IP allowlist. Otherwise, the IP allowlist does not take effect.

Add IP Blocklist or Allowlist

You can follow these steps to add an IP blocklist or allowlist:

1. Navigate to **Menu > System Management > IP Blocklist and Allowlist**.
2. Click **Add IP Blocklist and Allowlist**.

You can use the following example to complete the configuration:

- **Name:** The name of the IP blocklist or allowlist
- **Description:** The description of the IP blocklist or allowlist
- **Type:** Select blocklist or allowlist
- **IP Address:** You can enter IP addresses, IP address ranges, or IP/mask format. Separate multiple IP addresses with commas. You can add up to 100 entries.

Manage IP Blocklist and Allowlist

You can manage IP blocklists and allowlists, including editing names and descriptions, modifying configurations, and deleting them.

1. Navigate to **Menu > System Management > IP Blocklist and Allowlist**.
2. Select a list and then click **Action**.
 - 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 IP access restrictions for a particular list from the platform, select **Delete**.

13.2.2 Certificate Management

NexaSphere supports configuration and management of SSL certificates.

- [Import Third-Party Certificate](#)
- [Import System Self-Signed Certificate](#)
- [Update Certificate](#)
- [Switch to HTTP Login](#)

13.2.2.1 Import 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 > Certificate Management**.
2. On the **Certificate Management** page, click **Import Certificate**.
3. In the **Certificate Import** dialog, set the following parameters:
 - **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-----`.

- **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.

13.2.2.2 Import System Self-Signed 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.

Procedure

1. In the navigation pane, choose **System Management > Certificate Management**.

2. On the **Certificate Management** page, click **Import Certificate**.

3. In the **Certificate Import** dialog, set the following parameters:

- **Import Mode:** Select **System-Signed Certificate**.
- **Validity Period:** Choose from 3 months, 1 year, 3 years, 5 years, or 10 years. The default is 3 years.
- **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:

- **Common Name (CN):** Optional, set the common name; the default is localhost.

The length should be 1 to 64 characters, supporting only English letters in upper and lower case, numbers, and the following special characters: ~`@#%&*()- _ +={}[]!,:;<>./.

- **Organization (O):** Optional, set the organization name; the default is localhost.

The length should be 1 to 64 characters, supporting only English letters in upper and lower case, numbers, and the following special characters: ~`@#%&*()- _ +={}[]!,:;<>./.

- **Organizational Unit (OU):** Optional, set the department.

The length should be 1 to 64 characters, supporting only English letters in upper and lower case, numbers, and the following special characters: ~`@#%&*()- _ +={}[]!,:;<>./.

- **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, supporting Chinese characters, English letters in upper and lower case, numbers, and the following special characters: ~`@#%&*()- _ +={}[]!,:;<>./.

- **Locality (L):** Optional, set the city.

The length should be 1 to 128 characters, supporting Chinese characters, English letters in upper and lower case, numbers, and the following special characters: ~`@#%&*()- _ +={}[]!,:;<>./.

- **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.

13.2.2.3 Update 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.
- 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 > Certificate Management**.
2. In the **Certificate Management** 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.

13.2.2.4 Switch to HTTP Login

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 have configured the SSL certificate.

Procedure

1. In the navigation pane, choose **System Management > Certificate Management**.
2. In the **Certificate Management** 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.

13.2.3 Security Settings

NexaSphere provides security settings suitable for highly secure platform scenarios.

- [Modifying Security Settings](#)
- [Appendix: Security Settings Items](#)

Modifying Security Settings

You can follow these steps to modify security settings:

1. Navigate to the **Menu > System Management > Security Settings**.
2. Select the item you want to modify, then click the **Edit** icon to make changes.

Appendix: Security Settings Items

Category	Item Name	Item Description
Login Policy	Prohibit Multiple Session Connections for the Same User	Default is off. This setting determines whether multiple session connections for the same user are prohibited. If enabled, only one login session is allowed for the same user, and historical sessions will be forcibly terminated.
	Session Timeout	Default is 2 hours, measured in seconds/minutes/hours /days. After the session time exceeds this duration, the system becomes unavailable and requires re-login.
	Platform Login Verification Code Policy	Default is off. This setting determines whether the verification code function in login control is enabled. If enabled, after exceeding the maximum number of consecutive login failures, the verification code protection mechanism is triggered, requiring the correct username, password, and verification code to successfully log in to the platform. The default maximum number of consecutive login failures is 6 times.
	Platform Login Password Update Cycle	Default is off. This setting determines whether the password should be changed periodically. If enabled, when the password usage time reaches the specified update cycle, a prompt to change the password will

Category	Item Name	Item Description
		appear upon re-login, with the default being 90 days . When resetting the password, the new password cannot repeat any previously used passwords. The non-repetition count can be configured, with the default being 5, indicating that the new password cannot match any of the previous 3 used passwords.
	Platform Consecutive Login Failure User Lockout	Default is off. This setting determines whether consecutive login failures lock the user. If enabled, the account will be locked for a period of time after the specified number of consecutive login failures. The default maximum number of consecutive login failures is 6, and the default lockout duration is 10 minutes.
	Platform Login Password Strength	Default is off. If enabled, manual setting of password length and selection of whether to enable a combination of numbers, upper and lower case letters, and special characters is possible.
	Platform Login Two-Factor Authentication	Default is off. This setting determines whether two-factor authentication is enabled for logging into the platform.
Virtual Machine	VNC Console Password	Default is off. This setting determines whether password login to the VNC console is enabled. Note: The VNC password length range format is m-n, with values ranging from [6, 8] integers, defaulting to 6-8, and supporting the option to enable a combination of numbers, upper and lower case letters, and special characters.
	Virtual Machine Password Strength	Default is off. This setting determines whether password login to the virtual machine is enabled. 1. The virtual machine password length range format is m-n, with values ranging from [8, 32] integers, defaulting to 8-18, and supporting the option to enable a combination of numbers, upper and lower case letters, and special characters; 2. To set the virtual machine password, ensure that cloud-init is installed in the virtual machine image, and recommended versions of cloud-init are: 0.7.9, 17.1, 19.4, and versions after 19.4.
Host	Host Password Encryption Storage	Default is None. This setting determines the encryption storage policy for host passwords in the database. Available strategies are: None, LocalEncryption.

Category	Item Name	Item Description
		• None: No encryption storage. • LocalEncryption: Encryption storage using the built-in encryption feature of the platform.

13.3 System Settings

13.3.1 AccessKey Management

AccessKey is the security credential for authorizing third-party users to call the NexaSphere API to access virtualization resources. It includes the AccessKey ID and the AccessKey Secret, which must be kept strictly confidential.

- [Relationship Between AccessKey Administrator and User Permissions](#)
- [Generating an AccessKey](#)
- [Manage AccessKey](#)

Relationship Between AccessKey Administrator and User Permissions

- 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.

Generating an AccessKey

You can follow these steps to generate an AccessKey:

1. Navigate to **Menu > System Management > AccessKey Management**.
2. Click **Create AccessKey**.

Manage AccessKey

You can manage generated AccessKeys, including enabling, disabling, and deleting them.

1. Navigate to **Menu > System Management > AccessKey Management**.
2. Select the AccessKey.
 - To enable or disable an AccessKey, click **Disable** or **Enable**.
 - If you no longer need to access virtual resources via API, you can delete the AccessKey by clicking **Delete**.

13.3.2 Console Proxy Management

The console proxy enables logging into the virtual machine console through a proxy address. You can view, reconnect, and modify the console proxy address in NexaSphere.

- [View Proxy Address](#)
- [Reconnect Proxy Address](#)
- [Modify Proxy Address](#)

View Proxy Address

Procedure

1. Navigate to **Menu > System Management > Console Proxy**.
2. View the console proxy address. The default proxy address is the IP address of the management node.

Reconnect Proxy Address

If the virtual machine console fails to open, you need to perform a reconnect operation. After reconnecting, when the status displays as **Connected**, you can normally open the console. On the Console Proxy page, you can click **Reconnect** to perform the reconnect operation.

Modify Proxy Address

If you need to set the console proxy address, you can do so on the Console Proxy page by clicking **Modify Proxy Address**. This allows you to modify the address and port as needed. For the console proxy address, you can enter the public IP address of the management node, a NAT address, or a domain name. The changes take effect immediately after modification, without requiring a restart of the management node.

13.3.3 SNMP Management

13.3.3.1 Overview

You can monitor NexaSphere resource data and receive alert messages pushed by NexaSphere through the Simple Network Management Protocol (SNMP) on a third-party platform.

Key Concepts

- **SNMP Protocol:** A protocol used for managing devices on a network.
- **Network Management System:** A system that monitors and manages network devices via SNMP. This system sends requests to the agent process on managed devices to query

parameter values or receives Trap information sent proactively by the agent. 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 sends Trap information proactively. In this feature, the NexaSphere management node undertakes the role of the SNMP Agent.
- **Managed Object:** In this feature, it refers to resources on NexaSphere.
- **MIB Library:** A database maintained by the SNMP Agent that defines a set of attributes for managed objects, including object names, statuses, access permissions, and data types. The third-party platform collects required resource monitoring data based on instructions in the MIB library or parses received alert messages. NexaSphere provides a dedicated MIB library that supports downloading and viewing.
- **SNMP Trap Receiver:** A third-party server that receives alarm messages from NexaSphere. It can be added as an endpoint and attached to an alarm to push specified alarm messages.

Notes

The SNMP components has been updated at NexaSphere 1.10.25 version. If you have upgraded from a previous version and integrated with an external monitoring platform, you must:

1. Download the new MIB file.
2. Update the new MIB file to your external monitoring platform.



Note:

Failure to update may result in missing or abnormal monitoring data, affecting business monitoring continuity.

13.3.3.2 Enable SNMP Management

Procedure

1. In the navigation pane, choose **System Management > SNMP Management**.
2. On the **SNMP Management** page, click **Enable SNMP Management**.
3. In the **Enable SNMP Management** dialog, set the following parameters:

Basic Configuration

- **SNMP Agent Port:** Specify a port for receiving and responding to requests from the third-party monitoring platform. Default: 1160. Valid range: from 1024 to 65535.
- **Protocol Version:** Support v2c and v3 types.

- If you select the v2c type, set the **Community String** for connection authentication between the third-party monitoring platform and the virtualization platform.
- If you select the v3 type, set the following parameters:
 - **Username:** Set a username.
 - **User Authentication:** For secure authentication between the third-party monitoring platform and the platform, enable this option and set an authentication protocol and password.
 - **Data Encryption:** To encrypt communication messages between the virtualization platform and the third-party platform, enable this option and set an encryption protocol and password.

Receiver Configuration

- **SNMP Trap Receiver:** Enter the name, IP address, and port of the third-party server to receive pushed alert messages.

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

13.3.3.3 Modify SNMP Configuration

Procedure

1. In the navigation pane, choose **System Management > SNMP Management > SNMP Configurations**.
2. On the **SNMP Configurations** page, click **Modify Configuration**.
3. Make changes to the configuration as needed.

13.3.3.4 Download MIB File

Download the MIB file provided by the NexaSphere for data collection and alert message parsing on third-party platforms.

Procedure

1. In the navigation pane, choose **System Management > SNMP Management > SNMP Configurations**.
2. On the **SNMP Configurations** page, click **Download MIB**.

13.3.3.5 Manage SNMP Trap Receiver

Procedure

1. In the navigation pane, choose **System Management > SNMP Management > SNMP Trap Receiver**.
2. On the **SNMP Trap Receiver** page, follow these steps:
 - If you need to add a new receiver, click **Add SNMP Trap Receiver**.
 - If you need to modify the configuration information of an existing receiver, select the target object and then click **Actions > Modify Configuration**.
 - If you need to delete a receiver, select the target object and then click **Actions > Delete**.

**Note:**

If the selected SNMP Trap receiver has been added as an alert notification target, this notification target will also be deleted synchronously.

13.3.3.6 Disable SNMP Management

Procedure

1. In the navigation pane, choose **System Management > SNMP Management > SNMP Configurations**.
2. On the **SNMP Configurations** page, click **Disable**.

Result

- After disabling, the third-party platform will no longer actively obtain resource monitoring data from the platform.
- After disabling, the platform will continue to push alert messages to SNMP Trap receivers that have been added as notification targets.
- The platform will retain the current SNMP configuration for direct use the next time it is enabled.

13.3.4 Time Configuration

13.3.4.1 Overview

Manage platform time and configure time servers. After a time server is added, all host on the platform will synchronize time with the specified time server.

Definitions

The time configuration service involves the following concepts:

- Internal time server: Uses a management node or host as the time server for platform system time to synchronize time with other nodes on the platform.
- External time server: Uses an external node as the external time server to synchronize time with all nodes on the platform directly or via the internal time server.
- Time synchronization: The process in which the time of a node on the platform is synchronized with a time server.

Fundamentals

The local time is synchronized with the time server via network time protocol and certain algorithms. The following list shows the basic time synchronization process:

1. Select time source: Configure either internal or external time server as reference clocks.
2. Obtain timestamps: Each node on the platform communicates with the time server to acquire timestamps.
3. Calculate time difference: Based on the acquired timestamps, each node calculates the time difference between itself and the time server.
4. Adjust local clock: Each node synchronizes its local clock according to the time difference.
5. Periodic synchronization: Each node periodically communicates with the time server to prevent clock skew and maintains continuous time synchronization.

Benefits

The time configuration service has the following benefits:

- Precision: Uses accurate clock adjustment algorithms to achieve more precise time synchronization.
- Automation: Features an automatic time mechanism that periodically calibrates time without manual intervention.
- Reliability: Supports adding multiple time servers to enhance the reliability and stability of time synchronization.
- Intuitive: Displays the relationship between time servers and platform time configuration in a diagram and shows the current platform time and timezone.

Scenarios

The time configuration service is primarily used in the following scenarios:

- Network management: Provides precise time synchronization for analyzing log information collected from different network devices to facilitate fault localization.

- Billing system: Maintains uniform timekeeping to ensure accurate billing records.
- Collaborative processing: Ensures proper execution order when multiple systems process complex events simultaneously.

13.3.4.2 Modify Time Server Configuration

Context

The following NTP modes are supported:

- Internal: Uses a management node or host as the time server for the platform system time to synchronize time with other nodes on the platform. You can add a maximum of two internal time servers.
- External: Uses an external node as the external time server to synchronize time with all nodes on the platform. You can add a maximum of two external time servers.
- Internal and external: Uses an external node as the external time server and a management node or host as the internal time server. After the external time server synchronizes with the internal time server, the internal time server then synchronizes time with other nodes on the platform. You can add a maximum of two internal time servers and two external time servers.

Procedure

1. In the navigation pane, choose **System Management > Time Configuration**.
2. On the **Time Configuration** page, click **Modify Configuration**.
3. On the **Time Configuration** dialog, set the following parameters:
 - **NTP Mode**: Modify the NTP mode. Options include internal, external, and internal and external.
 - **Time Server**: add time servers based on the selected NTP mode. You can add a maximum of two internal time servers and two external time servers. External time servers can be added via IP or domain name.



Note:

- To modify a time server, make sure the management node is in connected status.
- Modifying the time server may cause platform node time inconsistencies, monitoring data deviations or errors, and impact running tasks. Carefully evaluate before making changes.

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

13.3.4.3 Synchronize Time

If the system time deviates significantly from the NTP time server, you can forcefully synchronize the time, avoiding the long duration of gradual adjustment.

Procedure

1. In the navigation pane, choose **System Management > Time Configuration**.
2. On the **Time Configuration** page, click **Sync Time**.

13.3.5 Log Server

13.3.5.1 Add Log Server

A log server can be used to collect management node logs, enabling quick issue identification and improving the operational efficiency of the platform.

Prerequisites

- You must have admin permissions to configure the log server.
- Make sure the communication between the management node and the log server is properly established.
- Make sure the log server has `syslog` server installed.

Procedure

1. In the navigation pane, choose **System Management > Log Server**.
2. On the **Log Server** page, click **Add Log Server**.
3. In the **Add Log Server** dialog, set the following parameters:
 - **Name**: Set the name of the log server.
 - **Description**: Optionally fill in a description for the log server.
 - **IP Address**: Enter the IP address of the log server.
 - **UDP Port**: Enter the port number that provides service for the UDP protocol.
 - **Log Identifier**: Select the log device category to match the log server, supporting LOCAL0~LOCAL7.

**Note:**

The level must be consistent with the setting in the log server's `rsyslog.conf` file to properly receive log information.

4. (Optional) Click **Test Connection** to check the IP address connectivity.
5. Review the configuration and click **OK**.

13.3.6 Email Server

13.3.6.1 Add Email Server

Procedure

1. In the navigation pane, choose **System Management > Email Server**.
2. On the **Email Server** page, click **Add Email Server**.
3. In the **Add Email Server** dialog, set the following parameters:
 - **Name**: Set the name of the email server.
 - **Description**: Optionally fill in a description for the email server.
 - **Username**: Enter the username for the email server.
 - **Password**: Enter the password corresponding to the username.
 - **Email Server Type**: The default is SMTP protocol.
 - **SMTP Server**: Enter the address of the email server.
 - **Encryption Type**: Choose whether to set up an encrypted connection for the email server port, including STARTTLS, SSL/TLS, or unencrypted.
 - **SMTP Port**: Set the port number for the email server.
4. Review the configuration and click **OK**.

13.3.6.2 Manage Email Server

Procedure

1. In the navigation pane, choose **System Management > Email Server**.
2. On the **Email Server** page, follow these steps:
 - If you need to modify the name and description of the email server, click **Actions > Edit Name and Description**.
 - If you need to enable or disable the email server, click **Actions** then **Enable/Disable**.
 - If you need to delete the email server, click **Actions > Delete**.

13.3.7 Theme Appearance

13.3.7.1 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 modify the theme appearance as needed.
3. For **Theme**, select the interface theme color. There are eight theme colors available for selection.
4. For **Browser Title**, set the following parameters:
 - **Favicon**: The browser icon only supports `.ico` format, with a file size not exceeding 2 MB.
 - **Chinese Title**: The Chinese title length for the browser should be within 25 characters.
 - **English Title**: The English title length for the browser should be within 25 characters.
5. For **Login Interface Title**, set the following parameters:
 - **Logo**: The login page logo image supports `.jpg`, `.jpeg`, `.png`, and `.svg` formats. The logo should be within 250×70 pixels, with a file size not exceeding 2 MB.
 - **Chinese Title**: The Chinese title length for the login page should be within 25 characters.
 - **English Title**: The English title length for the login page should be within 25 characters.
6. For **Platform Interface Title**, set the following parameters:
 - **Logo**: The platform interface logo image supports `.jpg`, `.jpeg`, `.png`, and `.svg` formats. The logo should be within 110×40 pixels, with a file size not exceeding 2 MB.

**Note:**

On dark backgrounds, it is recommended to use white or light-colored logos.

- **Chinese Title**: The Chinese title length for the platform interface should be within 25 characters.
- **English Title**: The English title length for the platform interface should be within 25 characters.
- **Font Size**: Select the text size for the platform interface title, including large, medium, and small. The default setting is medium.

13.3.7.2 Restore to Default Theme and Appearance

Restore the default theme appearance by clearing all current customizations with a single click.

Procedure

1. In the navigation pane, choose **System Management > Theme and Appearance**.
2. On the **Theme and Appearance** page, click **Restore Default Settings**.

13.3.8 System Parameters

NexaSphere provides system parameter functionality to globally control the default behavior of platform settings.

- [Overview](#)
- [Modify System Parameters](#)
- [Restore Default System Parameters](#)

Overview

The system parameters provided by the NexaSphere cover various aspects of platform configuration:

- Platform Policy: Offers system parameter settings at the platform level, including timeout policies, management node policies, deletion policies, cleanup policies, concurrency policies, reconnection policies, and progress bars.
- O&M Management: Provides system parameter settings related to monitoring alerts and task events.
- Host and VM: Offers system parameter settings related to hosts and virtual machines.
- Image Storage: Provides system parameter settings related to image storage and images.
- Data Storage: Offers system parameter settings related to data storage and disks.
- Network Resources: Provides system parameter settings related to distributed switches and security groups.

Modify System Parameters

To modify system parameters, follow these steps:

1. Navigate to **Menu > System Management > System Parameters**.
2. Set the system parameters as needed.

Restore Default System Parameters

To restore system parameters to their default state, follow these steps:

1. Navigate to **Menu > System Management > System Parameters**.
2. Click **Restore to Default Settings**.



Note:

Restoring default parameters will reset all system parameters, security settings, and advanced settings in high-availability policies to their initial default values. Proceed with caution.

14 Backup Management

14.1 Overview

The backup management module provides a set of functions for data protection and disaster recovery.

This section introduces the main features of backup management, as well as step-by-step guidance for data protection and disaster recovery.

Feature Overview

Centralized Management Interface

NexaSphere provides a centralized UI interface 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. Through backup policies, data retention policies, and backup QoS, you can customize and create backup plans that meet your business needs, and then apply these backup plans to virtual machines or platform databases that require backup protection.

Multiple Backup Policies

NexaSphere offers multiple backup policies for different backup objects, easily meeting virtualization backup requirements. For virtual machines, we provide default incremental backup and customized incremental backup policies. For the platform database, you can customize the backup cycle, execution time, and start time.

Flexible Multi-Scenario Backup

NexaSphere supports local and remote data protection, making data more reliable. Local backup storage supports seamless failover, effectively ensuring business continuity. Additionally, backup data can be synchronized from local backup storage to remote backup storage.

Backup Activity Monitoring

NexaSphere provides detailed information about backup-related activities. With just a few clicks, you can easily monitor the status of backup plan executions and the capacity usage of backup data.

14.2 Preparation

Before you use the backup feature in NexaSphere, make sure the platform version and license authorization meet the requirements.

- Make sure the installed software version is NexaSphere 4.3.0 or above.
- By default, NexaSphere provides a free trial quota for backing up 6 virtual machines. You do not need to purchase additional licenses to try out the backup feature.

Once the number of protected virtual machines reaches the free trial quota limit, if you wish to continue using the backup feature, purchase a Plus license. NexaSphere offers backup licenses in two ways: by the number of virtual machines and unlimited authorization. After secondary authorization, the authorized quota will override the free trial quota, for example: if the original free trial quota is 6 and the purchased authorized quota is 50, then the final effective authorized quota will be 50.

14.3 Add a Backup Storage

To use the backup feature, you need to deploy a backup storage in NexaSphere to store backup data. Depending on your backup scenario, you can add local backup storage or remote backup storage.

14.3.1 Add a Local Backup Storage

Add a local backup storage to store scheduled backup data for local virtual machines or the platform database. Local backup storage supports seamless failover to effectively ensure business continuity.

In case of accidental deletion or data corruption locally, you can restore backup data from local backup storage back to your local environment.

In the event of a disaster at the local data center, you can completely rely on the local backup storage to rebuild the data center and recover businesses.

You can choose any of the following methods to add local backup storage:

- Reuse an existing image storage to store backup data. For more information, see [Reuse Image Storage](#).
- Reuse an existing host and utilize free capacity on the host to store backup data. For more information, see [Reuse Host](#).

- Add a dedicated backup storage and use free disks or local directories on the storage to store backup data. For more information, see [Dedicated Backup Storage](#).

You can flexibly configure the number of local backup storage according to your actual business needs.

14.3.1.1 Reuse Image Storage

Reuse an existing image storage to store backup data.

Prerequisites

- The platform must have existing image storage resources, and these resources should be a standalone image storage that is connected and enabled.
- The backup network has been planned in advance.

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, set the following parameters:

Basic Information

- **Data Center**: Display the data center where the backup storage resides.
- **Addition Method**: Display **Reuse Image Storage**.

Configuration Information

- **Image Storage**: Select an existing image storage to be used as local backup storage.
- **Backup Storage Path**: Automatically retrieves the path of the selected image storage.
- **Backup Network**: The network dedicated for backups. Enter a backup network CIDR.
- **Backup Data**: Choose whether to scan for existing backup data.

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

14.3.1.2 Reuse Host

Reuse an existing host and utilize free capacity on the host to store backup data.

Prerequisites

- The platform must have existing host resources, and these hosts should be connected and enabled.
- The backup network has been planned in advance.

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, set the following parameters:

Basic Information

- **Name:** Name of the local backup storage.
- **Description:** Description of the local backup storage.
- **Data Center:** Display the data center where the backup storage resides.
- **Addition Method:** Display **Reuse Host**.

Configuration Information

- **Host:** Select an existing host to be used as local backup storage.
- **Backup Storage Path:** Enter the mount path on the local backup storage.



Note:

Avoid using system directories such as `/`, `/dev/`, `/proc/`, `/sys/`, `/usr/bin`, and `/bin`. Using system directories might cause hosts unable to work properly.

- **Backup Network:** The network dedicated for backups. Enter a backup network CIDR.
- **Backup Data:** Choose whether to scan for existing backup data.

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

14.3.1.3 Dedicated Backup Storage

Add a dedicated backup storage and use free disks or local directories on the storage 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 box, complete the basic configuration and disk configuration.

5. For **Basic Configuration**, set the following parameters:

Basic Information

- **Name:** Name of the local backup storage.
- **Description:** Description of the local backup storage.
- **Data Center:** Display the data center where the backup storage resides.
- **Addition Method:** Display **Dedicated Backup Storage**.

Configuration Information

- **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 corresponding to the user.

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

7. For **Disk Configuration**, set the following parameters:

- **Storage Method:** Supports **Free Disk** and **Local Directory**.
- **Free Disk:** When the storage method is set to **Free Disk**, this parameter must be configured.



Note:

This option will format the selected disks and completely erase all partitions, file systems, and data on the disk.

- **Backup Storage Path:** Enter the mount path on the dedicated backup storage.



Note:

Avoid using system directories such as `/`, `/dev/`, `/proc/`, `/sys/`, `/usr/bin`, `/bin`. Using system directories might cause backup storage unable to work properly.

- **Backup Network:** The network dedicated for backups. Enter a backup network CIDR.
- **Backup Data:** When the storage method is set to **Local Directory**, this parameter must be configured. Choose whether to scan for existing backup data.

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

14.3.2 Add Remote Backup Storage

Add a 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.

In case of accidental deletion or data corruption locally, you can restore backup data from remote backup storage back to your local environment.

In the event of a disaster at the local data center, you can completely rely on the remote backup storage to rebuild the data center and recover businesses.

You can add a remote backup storage by adding a dedicated backup node. For more information, see [Dedicated Backup Storage](#).

You can only add one remote backup storage.

14.3.2.1 Dedicated Backup Storage

Add a 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 box, complete the basic configuration and disk configuration.
5. For **Basic Configuration**, set the following parameters:

Basic Information

- **Name:** Name of the remote backup storage.
- **Description:** Description of the remote backup storage.
- **Data Center:** Display the data center where the backup storage resides.
- **Addition Method:** Display **Dedicated Backup Storage**.

Configuration Information

- **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 corresponding to the user.
6. Review the configuration and click **Next**.
7. For **Disk Configuration**, set the following parameters:
- **Storage Method:** Supports **Free Disk** and **Local Directory**.
 - **Free Disk:** When the storage method is set to **Free Disk**, this parameter must be configured.

**Note:**

This option will format the selected disks and completely erase all partitions, file systems, and data on the disk.

- **Backup Storage Path:** Enter the mount path on the dedicated backup storage.

**Note:**

Avoid using system directories such as `/`, `/dev/`, `/proc/`, `/sys/`, `/usr/bin/`, `/bin/`. Using system directories might cause backup storage unable to work properly.

- **Backup Network:** The network dedicated for backups. Enter a backup network CIDR.
 - **Backup Data:** When the storage method is set to **Local Directory**, this parameter must be configured. Choose whether to scan for existing backup data.
8. Review the configuration and click **OK**.

14.4 Virtual Machine Backup

In NexaSphere, the backup service is based on periodic scheduled backups. To perform a backup, you need to configure a backup plan. For more information, see [New Backup Plan](#).

During executions of backup plans, the backup mode, backup chain length, backup data retention policy, and backup QoS all depend on the configured backup policy. For more information, see [Backup Policy](#).

After you performed backups, you can use them to restore virtual machines. For more information, see [Data Recovery](#).

14.4.1 Backup Policy

Backup Policy

Backup Mode

NexaSphere provides two backup modes:

- **Default Incremental Backup**

You only need to customize the incremental backup policy, while the full backup policy is set by default. After 63 incremental backups, the system automatically performs a full backup, with a maximum backup chain length of 64.

- **Customized Incremental Backup**

You need to customize both the incremental and full backup policies. By customizing the full backup policy, you can shorten the interval between full backups. Increasing the number of full backups shortens the backup chain and reduces the amount of backup data that needs to be merged during recovery.

Backup Cycle

The backup cycle determines the frequency of backup execution.

- For the incremental backup policy, NexaSphere provides multiple options including monthly, weekly, daily, hourly, and minute-based backups.
- For the full backup policy, NexaSphere supports monthly and weekly backups.

Execution Time

You can set more granular backup execution times, down to the minute level. For the incremental backup policy, when the backup cycle is set to monthly, weekly, or daily, you can set the incremental backup execution time. For the full backup policy, when the backup cycle is set to monthly or weekly, you can set the full backup execution time.

- **Backup by Month:** If you set the execution time to 00:00 on the first day of each month, the incremental backup is executed at 00:00:00 on the first day of each month.
- **Backup by Week:** If you set the execution time to 00:00 on Sunday and Tuesday of each week, the incremental backup is executed at 00:00:00 on Sunday and Tuesday of each week.
- **Backup by Day:** If you set the execution time to 00:30, the incremental backup is executed at 00:30:00 every day.

Start Time

The backup plan begins at the start time and executes backups according to the configured backup cycle and execution time.

Retention Policy

Each successful backup generates a set of backup data used to restore virtual machine data to a previous point in time. To control the amount of backup data, you must specify a retention policy. Based on the backup scenario, NexaSphere can set separate retention policies for local backup data and remote backup data.

- Local Retention Policy:
 - Retain by quantity: Retain a minimum of 1 local incremental backup data and 1 local full backup data.
 - Retain by time: Retain a minimum of 1 day of local backup data.
- Remote Retention Policy:
 - Retain permanently: Remote backup data will not be automatically cleaned up.
 - Retain by quantity: Retain a minimum of 1 remote incremental backup data and 1 remote full backup data.
 - Retain by time: Retain a minimum of 1 day of remote backup data.



Note:

For data generated beyond the retention policy, only the backup records are deleted, which does not affect the data security.

14.4.2 New Backup Plan

Prerequisites

- Make sure the platform has sufficient authorized backup quota.
- Make sure the virtual machines to be backed up are in a running or paused state and are not associated with any other backup plans.
- You cannot associate a backup plan with a virtual machine, if it uses ZHPS distributed storage.

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**, set the following parameters:

Basic Information

- **Name:** Name of the backup plan.
- **Description:** Description of the backup plan.
- **Data Center:** Select a data center.

Backup Storage

- **Local Backup Storage:** Add local backup storage, with a maximum of 2 backup storage allowed for seamless failover.
- **Sync to Remote Backup Storage:** Disabled by default. You can choose whether to synchronize local backup data to a remote backup storage.



Note:

If enabled, ensure that remote backup storage has been added in advance on the platform. For more information, see [Add Remote Backup Storage](#).

- **Remote Backup Storage:** After enabling synchronization to remote backup storage, it displays the name and available capacity of the remote backup storage.

Backup Object

- **Backup Object Type:** Select **Virtual Machine**.
- **Virtual Machine:** Add virtual machines that need 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.
- If the selected virtual machine already has backup data, creating a new backup plan will not additionally consume authorized backup quota.
- By default, the entire VM is backed up. However, shared disks and RDM disks attached to VMs are excluded from backup.

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

6. For **Backup Configuration**, set the following parameters:

Backup Policy

- **Backup Mode:** Select default incremental backup or customized 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 set to customized incremental backup, you need to customize the full backup policy, including the backup cycle and execution time.
- **Maximum Backup Chain Length:** Display the length of the backup chain under the current policy, with a maximum length not exceeding 64.
- **Start Time:** Set a start time for the backup plan.
- **Disk Read Speed:** Set a maximum disk read speed limit. Default value: unlimited. Units: MB/s, GB/s. Range: 1MB/s to 100GB/s.

**Note:**

QoS settings should match the physical network bandwidth and account for the bandwidth occupied by concurrent backups.

For more information about backup policies, see [Backup Policy](#).

Retention Policy

- **Local Retention Policy:** Set the local backup retention policy by quantity or time.
- **Remote Retention Policy:** After enabling **Sync to Remote Backup Storage**, this parameter must be set. You can set the remote retention policy as permanently or by quantity or time.

For more information about retention policies, see [Backup Policy](#).

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

What's next

After you performed backups, you can use them to restore virtual machines. For more information, see [Data Recovery](#).

14.4.3 View Backup Plan and Data

During executions of backup plans, you can monitor the progress and get real-time updates on the backup activities. For more information, see [View Backup Plan Executions](#).

During executions of backup plans, a series of backup data is generated, and you can view the backup data for protected resources. For more information, see [View VM Backup Data](#).

14.4.3.1 View Backup Plan Executions

View the Latest Backup Result

Procedure :

1. In the navigation pane, choose **Data Protection > Backup Plan**.
2. You can quickly check the results of the most recent backup through the **Last Backup Result** item in the list.

Check Backup Execution Status of Protected Resources

If you have initiated a large backup schedule involving multiple virtual machines, you can view the backup results for each virtual machine to understand which VM backups succeeded, failed, or are pending.

Procedure :

1. In the navigation pane, choose **Data Protection > Backup Plan**.
2. On the **Backup Plan** page, click the name of the backup plan to enter its details page.
3. On the backup plan details page, click **Backup Resources**.
4. On the **Backup Resources** tab, you can view the backup execution status of all protected resources in this backup plan.

View Backup Job Details in Backup Plan

You can view detailed information for each backup job, including the backup object, backup mode , number of backup resources, total backup size, backup job result, backup duration, start and completion times, as well as backup details.

Procedure :

1. In the navigation pane, choose **Data Protection > Backup Plan**.
2. On the **Backup Plan** page, click the name of the virtual machine backup plan to enter its details page.
3. On the backup plan details page, click **Backup Job**.
4. On the **Backup Job** tab, select the backup job you want to view and then click **View Backup Details**.
5. In the **Backup Job Details** dialog, you can view the detailed information about the backup job.

14.4.3.2 View VM Backup Data

During executions of backup plans, a series of backup data is generated and saved in the specified backup storage. NexaSphere provides a unified interface to help you view the backup data of protected resources, ensuring that no important business backups are overlooked.

Procedure

1. In the navigation pane, choose **Data Protection > Protected Resources**.
2. Navigate to the **Protected Resources** page and select the **Virtual Machine** tab.
3. View the backup data of virtual machines.

The virtual machine backup data page displays the backup data for each protected virtual machine in a tree-like folder structure, organized by individual VM. Under each virtual machine folder, all related backup data is listed from top to bottom based on the most recent creation date.

- After selecting a virtual machine, you can view its power status, total backup data size, associated backup plan, UUID, and all backup data for that virtual machine.
- After selecting a specific backup data of a virtual machine, you can view basic information and backup details about this data, including the associated virtual machine, backup size, backup type, associated backup storage, whether it has been synchronized to remote backup storage, owner, UUID, creation time, and detailed backup information.

14.4.4 Perform Backups Manually

14.4.4.1 Immediate Backup

If you wish to create an additional backup restore point for one or more virtual machines in a backup plan without creating a new backup plan or modifying the existing one, or if you have already associated a backup plan with the virtual machines you want to protect but no backup data has been generated according to the scheduled time, you can use manual backups.

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 perform a backup, then click **Actions > Backup Now**.
3. In the **Backup Now** dialog, choose whether to enable **Full Backup** as needed.



Note:

If enabled, it will manually perform a full backup once for the backup objects. If disabled, the backup mode will be the same as the next scheduled backup mode.

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

If your backup plan includes multiple virtual machines and you only want to generate backup data for a specific VM, you can first navigate to the **Backup Resources** tab of the backup plan details page. There, select the virtual machines that you do not wish to back up at this time, then click **Actions > Backup Job State > Disable**. Once successfully disabled, you can manually perform the backup for the backup plan, which will then only affect the active VMs.

14.4.4.2 Create On-Demand Backup

In addition to performing backups by adding virtual machines to a backup job, you can create an on-demand backup for a virtual machine.

Prerequisites

- You cannot create backups for a virtual machine, if it uses ZHPS distributed storage.
- Shared disks or RDM disks attached to the virtual machine will not be included in the backup.

Procedure

1. In the navigation pane, choose **Inventory > VM and Host**.
2. In the left navigation tree, right-click a virtual machine, then select **Backup > Create Backup**.
3. In the **Create Backup** dialog, enter a backup name, select the backup type, specify the local backup storage, and choose whether to synchronize to remote backup storage.
4. Review the configuration and click **OK**.

What's next

After you successfully created the backup, you can go to the **Protected Resources** page to view the backup data for that virtual machine.

14.5 Platform Database Backup

In NexaSphere, the backup service is based on periodic scheduled backups. To perform a backup, you need to configure a backup plan. For more information, see [New Backup Plan](#).

During executions of backup plans, the backup cycle, execution time, start time, and backup data retention policy all depend on the configured backup policy. For more information, see [Backup Policy](#).

After you performed backups, you can use them to restore the platform database. For more information, see [Data Recovery](#).

14.5.1 Backup Policy

Backup Policy

Backup Cycle

The backup cycle determines the frequency of backup execution. NexaSphere provides three options for platform database backups: weekly, daily, and hourly backups.

Execution Time

If the backup cycle is set to weekly or daily, you can set a more granular execution time, down to the minute level.

- Backup by Week: If you set the execution time to 00:00 on Sunday and Tuesday of each week, the incremental backup is executed at 00:00:00 on Sunday and Tuesday of each week.
- Backup by Day: If you set the execution time to 00:30, the incremental backup is executed at 00:30:00 every day.

Start Time

The backup plan begins at the start time and executes backups according to the configured backup cycle and execution time.

Retention Policy

Each successful backup generates a set of backup data used to restore the platform database data to a previous point in time. To control the amount of backup data, you must specify a retention policy. Based on the backup scenario, NexaSphere can set separate retention policies for local backup data and remote backup data.

- Local Retention Policy:
 - Retain by quantity: Retain a minimum of 1 local backup data.
 - Retain by time: Retain a minimum of 1 day of local backup data.
- Remote Retention Policy:
 - Retain permanently: Remote backup data will not be automatically cleaned up.
 - Retain by quantity: Retain a minimum of 1 remote backup data.
 - Retain by time: Retain a minimum of 1 day of remote backup data.

**Note:**

For data generated beyond the retention policy, only the backup records are deleted, which does not affect the data security.

14.5.2 New Backup Plan

Prerequisites

- Make sure the platform has sufficient authorized backup quota.
- NexaSphere supports only one platform database backup plan. Make sure that there is no existing platform database 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**, set the following parameters:

Basic Information

- **Name:** Name of the backup plan.
- **Description:** Description of the backup plan.
- **Data Center:** Select a data center.

Backup Storage

- **Local Backup Storage:** Add local backup storage, with a maximum of 2 backup storage allowed for seamless failover.
- **Sync to Remote Backup Storage:** Disabled by default. You can choose whether to synchronize local backup data to a remote backup storage.

**Note:**

If enabled, ensure that remote backup storage has been added in advance on the platform. For more information, see [Add Remote Backup Storage](#).

- **Remote Backup Storage:** After enabling synchronization to remote backup storage, it displays the name and available capacity of the remote backup storage.

Backup Object

- **Backup Object Type:** Select Platform Database.

**Note:**

If a platform database backup plan already exists, the platform database option will not be displayed.

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

6. For **Backup Configuration**, set the following parameters:

Backup Policy

- **Backup Cycle:** Select the backup cycle as needed, including weekly, daily, or hourly backups.
- **Execution Time:** After selecting a weekly or daily backup cycle, you need to set the execution time.
- **Start Time:** Set a start time for the backup plan.

For more information about backup policies, see [Backup Policy](#).

Retention Policy

- **Local Retention Policy:** Set the local backup retention policy by quantity or time.
- **Remote Retention Policy:** After enabling **Sync to Remote Backup Storage**, this parameter must be set. You can set the remote retention policy as permanently or by quantity or time.

For more information about retention policies, see [Backup Policy](#).

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

What's next

After you performed backups, you can use them to restore the platform database. For more information, see [Data Recovery](#).

14.5.3 View Backup Plan and Data

During executions of backup plans, you can monitor the progress and get real-time updates on the backup activities. For more information, see [View Backup Plan Executions](#).

During executions of backup plans, a series of backup data is generated, and you can view the backup data for protected resources. For more information, see [View Platform Database Backup Data](#).

14.5.3.1 View Backup Plan Executions

View the Latest Backup Result

Procedure:

1. In the navigation pane, choose **Data Protection > Backup Plan**.
2. You can quickly check the results of the most recent backup through the **Last Backup Result** item in the list.

View Backup Job Details in Backup Plan

You can view detailed information for each backup job, including the backup object, backup mode, number of backup resources, total backup size, backup job result, backup duration, start and completion times, as well as backup details.

Procedure

1. In the navigation pane, choose **Data Protection > Backup Plan**.
2. On the **Backup Plan** page, click the name of the virtual machine backup plan to enter its details page.
3. On the backup plan details page, click **Backup Job**.
4. On the **Backup Job** tab, select the backup job you want to view and then click **View Backup Details**.
5. In the **Backup Job Details** dialog, you can view the detailed information about the backup job.

14.5.3.2 View Platform Database Backup Data

NexaSphere provides a unified interface to help you view the backup data of protected resources, ensuring that no important business backups are overlooked.

Procedure

1. In the navigation pane, choose **Data Protection > Protected Resources**.
2. Navigate to the **Protected Resources** page and select the **Platform Database** tab.
3. View the backup data of the platform database.

The platform database backup data page is divided into two tabs based on the data storage location:

- Local Backup: Backup files are stored in local backup storage.
- Remote Backup: Backup files are stored in remote backup storage.

After you selected the corresponding location, you can view the name of the platform database backup, the management node IP, version, capacity, backup storage path (or whether it has been synchronized to remote/local storage), and creation time.

14.6 Data Recovery

NexaSphere provides multiple data recovery methods to address various disaster recovery scenarios:

- You can restore an entire virtual machine from backup data to a specified location. For more information, see [Restore Virtual Machine](#).
- You can create a new virtual machine directly from the backup data. For more information, see [New Virtual Machine from Backup](#).
- You can restore the platform from backup data to a specified location. For more information, see [Restore Platform Database](#).

14.6.1 Restore Virtual Machine

If a virtual machine fails, you can restore the virtual machine from backup data to the latest backup time point or to a previous backup time point.

Prerequisites

- Make sure the virtual machine has at least one successfully generated backup data.
- To restore a virtual machine, you need to power off the virtual machine.
- To use remote backup data to restore a virtual machine, you need to synchronize the remote backup data to local backup storage before proceeding with the recovery.

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 you want to restore, then click **Restore**.
4. In the **Restore Virtual Machine** dialog, you can choose whether to automatically power on after the recovery as needed.
5. Review the configuration and click **OK**.



Note:

Restoring a VM overwrites the existing data with the selected backup data. Proceed with caution.

14.6.2 New Virtual Machine from Backup

If a virtual machine fails, you can create a new virtual machine directly from the backup data.

This data recovery method is non-destructive and does not overwrite the existing virtual machine.

Instead, it creates a new virtual machine based on the backup data.

Prerequisites

- Make sure the virtual machine has at least one successfully generated backup data.
- Make sure the platform has sufficient compute, storage, and network resources to support the new virtual machine.
- To use remote backup data to create a new virtual machine, you need to synchronize the remote backup data to local backup storage before proceeding with the creation.

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 you want to restore, then click **New Virtual Machine**.
4. In the **New Virtual Machine from Backup** dialog, set the following parameters:

Backup Information

- **Backup Data:** Displays the selected backup data.

Basic Information

- **Name:** Name of the virtual machine.
- **Quantity:** Default is 1, modification is not supported.
- **Group:** Group where the virtual machine resides. If not set, the default group will be used.
- **Location:** Host or cluster location where the virtual machine resides.
- **OS:** Display the operating system of the virtual machine recorded in the backup data.
- **HA:** Automatic restart mechanism after an abnormal shutdown of the virtual machine. For more information, see [VM HA](#).
- **Power Status:** Whether to automatically power on the virtual machine after the creation.

Hardware Information

- **CPU:** Support adjusting the total number of cores.
- **Memory:** Support adjusting the memory size.
- **Disk:** Display the disk configuration recorded in the backup data. Modification is not supported.
- **NIC:** Support adjusting port groups, MAC address, IP address, DNS assignment, and security groups

You can add a new NIC to the virtual machine by clicking **Add NIC**. The new NIC allows customization of the network address and features.

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

What's next

Some VM configurations require VMTools. After VM creation, it is recommended to install VMTools to enable certain configurations. For more information about VMTools, see [Virtual Machine VMTools](#).

14.6.3 Restore Platform Database

Restore the platform from backup data to the state at a specified backup time point.

Prerequisites

- Restoring the platform database requires a management node restart, during which the management interface will be unavailable. This process typically takes several minutes, and your business resources will not be affected.
- Resources that were deleted after the backup point will become invalid data and cannot be restored to normal during the recovery of 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 files, such as **Local Backup** or **Remote Backup**.
4. In the platform database backup list, select the backup data you want to restore, then click **Actions > Restore**.
5. In the **Restore Platform Database** dialog, enter the platform database root password.
6. Review the configuration and click **OK**.

What's next

After successfully restored the database, all platform resources will revert to their state at the time of backup creation. Navigate to the **Data Protection > Backup Storage** page, select the backup storage, and then click **Actions > Scan Backup Data** to obtain current and accurate backup data information.

14.7 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 to remote backup storage, or synchronize one or more remote backup data to local backup storage, providing dual protection for your backup data. This section uses virtual machine backup data as an example to introduce how to synchronize local backup data to remote backup storage.

Prerequisites

- The remote backup storage must be added to the platform in advance. For more information, see [Add Remote Backup Storage](#).
- The selected backup data to be synchronized has not been previously synchronized.

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 virtual machine whose backup data you want to synchronize, then in the right-side **Backup Data** list, check the backup data you want to synchronize.
5. Click **Bulk Actions > Sync to Remote Backup Storage**.
6. In the **Sync to Remote Backup Storage** dialog , specify a remote backup storage.
7. Review the configuration and click **OK**.

Change the Owner of VM Backup Data

You can change the owner of virtual machine 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 virtual machine, then in the right-side **Backup Data** list, check the backup data for which you want to change the owner.
4. Click **Bulk Actions > Change Owner**.
5. In the **Change Owner** dialog, specify the new owner.
6. Review the configuration and click **OK**.

Export Platform Database Backup Data

You can export the backup data of the platform database 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 files, such as **Local Backup** or **Remote Backup**.
4. In the platform database backup list, select the platform database backup you want to export, then click **Actions > Export**.

Delete Backup Data

After you have determined that the backup data of protected resources is no longer needed, you can delete the backup data. This section uses virtual machine backup data as an example to illustrate how to delete 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 virtual machine, then in the right-side **Backup Data** list, check the backup data you want to delete.
4. Click **Bulk Actions > Delete**.

5. In the **Delete Backup Data?** dialog, you can choose whether to delete remote backup data at the same time.
6. Review the configuration and click **OK**.

14.8 Backup Plan Management

You can perform the following operations on backup plan:

- [Enable/Disable Backup Plan](#)
- [Edit Name and Description](#)
- [Modify Basic Settings](#)
- [Modify Backup Policy](#)
- [Delete Backup Plan](#)

Enable/Disable Backup Plan

You can temporarily disable a regularly executed backup plan. A disabled backup plan is not deleted. Instead, it will be paused for a period of time and will not run according to the specified backup cycle. The associated protected resources will also stop the scheduled backup until you enable the backup plan again.

Procedure

1. In the navigation pane, choose **Data Protection > Backup Plan**.
2. On the **Backup Plan** page, select the backup plan you want to disable or enable, 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, then click **Actions > Edit Name and Description**.
3. In the **Edit Name and Description** dialog, enter the new name or description, then click **OK**.

Modify Basic Settings

You can modify the basic configuration of a backup plan or add more virtual machines to the 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, then click **Actions > Modify Basic Settings**.
3. In the **Modify Basic Settings** dialog, modify the basic information, backup storage information, and backup object information as needed.
4. Review the configuration and click **OK**.

Modify Backup Policy

You can modify the backup policy of a backup plan at any time, such as the backup policy, retention policy, and advanced settings.

For more information about virtual machine 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, then click **Actions > Modify Backup Policy**.
3. In the **Modify Backup Policy** dialog, modify the backup policy, retention policy, and advanced settings as needed.
4. Review the configuration and click **OK**.

Delete Backup Plan

If you determine that a backup plan is no longer needed, you can delete the backup plan. After deletion, associated resources will no longer perform scheduled backup actions. If you still need to back up specific resources, you can 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, then click **Actions > Delete**.

14.9 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 flexibly manage the status of backup storage, including enabling, disabling, and reconnecting.

Disabling backup storage may affect the execution of backup plans. Note that:

- If a backup plan specifies only one backup storage, disabling that backup storage will cause the backup plan to fail.
- If a backup plan specifies two backup storages, disabling one of the backup storages (primary) will cause the backup plan to automatically switch to the other backup storage (secondary) for backup execution.

Procedure

1. In the navigation pane, choose **Data Protection > Backup Storage**.
2. On the **Backup Storage** page, select the target data storage, then click **Actions > Enable/Disable/Reconnect**.

Scan Backup Data

You can scan the backup data in the backup storage that has not been deleted. Deleted backup data cannot be recovered through scanning. Alternatively, if you have manually added backup data to the backup storage and wish to use this 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 data storage, then click **Actions > Scan Backup Data**.

Clean Up Backup Storage Data

You can clean up invalid backup data that has been completely deleted from the backup storage to free up storage space.

Procedure

1. In the navigation pane, choose **Data Protection > Backup Storage**.
2. On the **Backup Storage** page, select the target data storage, then click **Actions > Cleanup Data**.

Edit Name and Description

You can modify the name or description of a 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 data storage, then click **Actions > Edit Name and Description**.
3. In the **Edit Name and Description** dialog, enter the new name or description, then click **OK**.

Modify Backup Storage Configuration

You can modify the configuration information of an added backup storage.

Procedure

1. In the navigation pane, choose **Data Protection > Backup Storage**.
2. On the **Backup Storage** page, select the target data storage, then click **Actions > Modify Configuration**.
3. In the **Modify Configuration** dialog, modify the basic information or configuration information as needed.
4. Review the configuration and click **OK**.

Update Backup Storage Password

You can modify the password of an added backup storage.

Procedure

1. In the navigation pane, choose **Data Protection > Backup Storage**.
2. On the **Backup Storage** page, select the target data storage, then click **Actions > Update Password**.
3. In the **Update Password** dialog, enter the new password.
4. Review the configuration and click **OK**.

Delete Backup Storage

Deleting backup storage is a high-risk operation that will stop all related backup plans on the backup storage, and related backup data will no longer be accessible.

Procedure

1. In the navigation pane, choose **Data Protection > Backup Storage**.
2. On the **Backup Storage** page, select the target data storage, then click **Actions > Delete**.

15 Bare Metal Management

15.1 Overview

NexaSphere Bare Metal Management Service provides dedicated physical servers for your critical applications, delivering high performance and stability.

After completing server installation and basic setup, you can rapidly deploy bare metal chassis and create bare metal instances on the UI. The service supports unattended OS installation through bare metal templates, significantly improving operational efficiency. You can flexibly configure business networks for bare metal instances to meet different application requirements.

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

The bare metal management service offers the following advantages:

- **Dedicated physical resources**

Provides dedicated physical servers for your critical applications, ensuring high performance and stability.

- **Efficient batch deployment**

- You can add bare metal chassis in batches on the UI through manual addition or template import. You can batch configure IPMI addresses for multiple chassis at once, quickly building bare metal clusters.
- You can use bare metal templates for unattended OS installation, significantly improving operational efficiency.

- **Flexible OS support**

Compatible with mainstream Linux distributions (RHEL/CentOS series, Debian/Ubuntu series , and SUSE/openSUSE series) and custom platform OS versions, meeting you customized installation requirements.

- **High-availability deployment architecture**

Supports dual management node HA scenarios. It is recommended to attach separate deployment server for each bare metal cluster to avoid single points of failure and DHCP conflicts.

Scenarios

The bare metal management service is suitable for the following scenarios:

- **High-security regulated environments**

The financial and insurance industries that have high requirements over business deployment compliance and data security. In these scenarios, you can use Bare Metal Management to secure dedicated resources, data isolation, easy management, and operation-tracking. This way, you can ensure the reliability and security compliance of your key business system and data.

- **High-performance computing**

In supercomputing, genome sequencing, and other high-performance computing scenarios, the requirements over the computing performance, stability, and timeliness of the server are very high. The Bare Metal Management feature is fitting for these scenarios. In addition, the feature can be used for scenarios that require high throughput or high computing performance that can accommodate changing access requests and scenarios. Virtualization and hyperthreading may compromise some performance. Deploying a reasonable number of bare metal clusters can meet the high-performance computing requirements.

- **Mission-critical databases**

To meet business requirements, you may not want to deploy some key databases on virtual machines while want to deploy the databases on physical servers that feature dedicated resources, network isolation, and guaranteed performance. In these scenarios, you can use Bare Metal Management to provide dedicated high-performance physical servers for your applications.

15.2 Preparation

License

To use the NexaSphere Bare Metal Management service, you need to purchase an individual plus license.

Network Planning

Refer to [Figure 15-1: Bare Metal Network Topology](#) to plan your physical network environment in advance, including the deployment network, IPMI network, management network, and business network.

Figure 15-1: Bare Metal Network Topology

1. Plan the deployment network.

Make sure the PXE NIC of bare metal chassis connects to the deployment server's (PXE server) DHCP listening NIC through the deployment network.

- Access the bare metal chassis BIOS to confirm the NIC connecting to the deployment network has PXE enabled.

For some models, make sure the PXE NIC is the first boot NIC or disable PXE on all NICs with higher boot priority.

- No manual BIOS boot order adjustment is needed (typically HDD boot first). The system will automatically trigger a one-time PXE boot.
- For high availability, we recommend using independent deployment servers.
- The deployment server's DHCP listening NIC must be a dedicated NIC with IP address to provide stable DHCP service.
- Ensure no other DHCP services exist in the deployment network.
- Plan IP address allocation for each bare metal instance according to your production environment.


2. Configure bare metal chassis IPMI and plan IPMI network.

Two configuration scenarios exist:

- For out-of-band control, ensure each bare metal chassis has BMC interface with pre-configured IPMI address, port, username and password. Plan IPMI network to connect management nodes to BMC interfaces.

- If devices support BMC ShareLink Management Port, disable it in BIOS to ensure only one IPMI address exists.

Figure 15-2: Disable BMC ShareLink Management Port



3. Plan management and business networks.

- Plan management network to connect image storage and deployment servers to management nodes.
- For bare metal instance business networks, attach distributed switches to the corresponding bare metal cluster in advance.

Software Preparation

Prepare ISO images in the image storage for installing operating systems on bare metal instances.

- The bare metal instance supports both custom platform OS and mainstream Linux distributions , including RHEL/CentOS series, Debian/Ubuntu series, and SUSE/openSUSE series.
- Only ISO format images are supported.

15.3 Quick Start Guide

Follow this workflow to quickly create bare metal instances:

1. (Optional) Add custom bare metal templates.

Use templates to generate pre-configured files for unattended batch OS installation on bare metal instances. You can use either system templates or custom templates. For more information, see [Bare Metal Template](#).

2. Create a new bare metal cluster.

Provides dedicated cluster management for bare metal devices. For more information, see [Bare Metal Cluster](#).

3. Add bare metal chassis.

You can add bare metal chassis by manual addition or template import. For more information, see [Bare Metal Chassis](#).

4. Create bare metal instances.

For more information, see [Bare Metal Instance](#).

15.4 Bare Metal Template

15.4.1 Template Syntax Rules

Bare metal templates include both system variables and custom variables to support various unattended deployment scenarios.

System templates only contain predefined system variables. Custom templates include both system variables and your custom variables.

System variables example (All uppercase, underscore-separated):

```
REPO_URL
# Installation source URL created from selected ISO
# Can be commented out and manually specified via --url

USERNAME
# System username
# Default is root for RHEL/CentOS or SUSE/openSUSE (only password
  required)
# Required for Debian/Ubuntu systems

PASSWORD
# Password for the system user

NETWORK_CFGS
# NIC UUID and network UUID from UI
# Replaced after IP assignment (automatic or manual)

FORCE_INSTALL
# Whether to overwrite existing disk data automatically

PRE_SCRIPTS
```

```
# Pre-installation scripts

POST_SCRIPTS
# Post-installation scripts
```

Custom variables example (All lowercase, underscore-separated):

```
hostname
# Hostname

keyboard
# Keyboard

timezone
# Timezone
```

Different template types follow different syntax rules:

- For kickstart templates: Refer to Red Hat official documentation.
- For preseed or autoinstall templates: Refer to Ubuntu official documentation.
- For autoyast templates: Refer to SUSE official documentation.

15.4.2 System Preconfigured Templates

NexaSphere provides multiple system templates.

Template Name	Description
cloud_host_x86_64_v3	For unattended deployment of NexaSphere management node via the bare metal module.
cloud_expert_x86_64_v2	For unattended deployment of NexaSphere expert mode.
centos_7_x86_64_mini_v1	For unattended deployment of CentOS 7 systems via the bare metal module.
centos_7_aarch64_min_v1	
kylin_10_x86_64_min_v1	For unattended deployment of Kylin V10 systems via the bare metal module.
kylin_10_aarch64_min_v1	
openEuler_20_aarch64_min_v1	For unattended deployment of openEuler 20 systems via the bare metal module.
opensuse_15_x86_64_mini_v1	For unattended deployment of openSUSE 15 systems via the bare metal module.
ubuntu_16_x86_64_mini_v2	For unattended deployment of Ubuntu 16 systems via the bare metal module.
ubuntu_18_x86_64_mini_v1	For unattended deployment of Ubuntu 18 systems via the bare metal module.

Template Name	Description
ubuntu_20_live_server_x86_64_mini_v1	For unattended deployment of Ubuntu 20 and Ubuntu 22 systems via the bare metal module.

15.4.3 Add a Custom Template

Upload UTF-8 encoded custom bare metal template files for complex unattended deployment scenarios.

Prerequisites

- Custom template file size must not exceed 50 KB.
- Custom template files must strictly follow syntax rules for their 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, set the following parameters:

- **Name:** Set a name for the bare metal template.
- **Description:** Optional. Enter a description.
- **Operating System:** Select the operating system for unattended deployment.

Supports both custom platform OS and mainstream Linux distributions, including RHEL/CentOS series, Debian/Ubuntu series, and SUSE/openSUSE series.

- **Template Type:** Select template type matching the OS.
 - Custom platform-optimized OS: Select kickstart.
 - RHEL/CentOS series: Select kickstart.
 - Debian/Ubuntu series: Select preseed.
 - Ubuntu Live: Select autoinstall.
 - SUSE/openSUSE series: Select autoyast.
- **Template Import:** Upload an UTF-8 encoded custom template file.



Note:

Custom bare metal templates must strictly follow the syntax rules of the corresponding selected template type.

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

15.5 Bare Metal Cluster

15.5.1 New Bare Metal Cluster

Prerequisites

- The bare metal cluster requires a deployment server to provide PXE services for bare metal instances in the cluster.
- Each bare metal cluster supports only one deployment server, but one deployment server can be attached to multiple clusters.
- The bare metal cluster can attach distributed switches to provide networks for bare metal instances.

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, click **Bare Metal Cluster > New Bare Metal Cluster**.
5. In the **New Bare Metal Cluster** dialog, set the following parameters:
 - **Name**: Set a name for the bare metal cluster.
 - **Description**: Optional. Enter a description.
 - **Attach Deployment Server**: Choose whether to attach a deployment server.

- **Server Source**: Select either a new or existing deployment server.

If you select **New Deployment Server**, set the following parameters:

- **Deployment Server Name**: Set a name for the deployment server.
- **DHCP Listening NIC**: Enter the NIC device number connected to the deployment network.



Note:

- This NIC must be connected to the deployment network of the bare metal chassis and have a configured IP address.

- The network where this NIC resides must not have any other DHCP services.
- **Storage Path:** Enter the local directory for PXE deployment images.
- **Deployment Server IP:** Enter the deployment server IP address.
- **SSH Port:** Default: 22.
- **Username:** Default: root. You can enter a regular user.
- **Password:** Enter corresponding user password.
- **DHCP Start IP:** Used for traversing IP ranges of DHCP services.
- **DHCP End IP:** Used for traversing IP ranges of DHCP services.

**Note:**

If not specified, the system will detect and filter used IP address as IP ranges according to this NIC IP addresses.

If you select **Existing Deployment Server**, select a deployment server.

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

15.5.2 Deployment Server Management

Procedure

1. In the navigation pane, choose **Inventory > Bare Metal Management**.
2. Select the target bare metal cluster.
3. On the bare metal cluster **Overview** details page, you can check the configuration information of the attached deployment server.
4. On the bare metal cluster details page, click **Actions**.
 - To attach a deployment server, choose **Attach Deployment Server**.
 - To reconnect a deployment server, choose **Reconnect Deployment Server**.
 - To detach the deployment server from the bare metal cluster, choose **Detach Deployment server**.

15.5.3 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, click **Actions > Delete**.

Result

Deleting bare metal clusters will delete all bare metal chassis and bare metal instances in the cluster. Proceed with caution.

15.6 Bare Metal Chassis

15.6.1 Add a Bare Metal Chassis

15.6.1.1 Manually Add a Bare Metal Chassis

Add single or multiple bare metal chassis by specifying an IPMI address or IPMI range.

Prerequisites

Make sure 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 choose **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, set the following parameters:
 - **Name**: Set a name for the bare metal chassis.
 - **Description**: Optional. Enter a description.
 - **Bare Metal Cluster**: Select the target bare metal cluster where the bare metal chassis reside.
 - **Addition Method**: Choose a method to add IPMI addresses. Options include IPMI address and IPMI range.
 - **IPMI Address/Range**: Enter an IPMI address or the start and end IP of an IPMI range.
 - **Port**: Enter the IPMI port. Default: 623.
 - **Username**: Enter the IPMI username.
 - **Password**: Enter the IPMI password.
 - **Reboot Bare Metal Chassis**: Choose whether to reboot the bare metal chassis to automatically obtain the hardware information. Default: unselected.
6. Review the configuration and click **OK**.

15.6.1.2 Import a Bare Metal Chassis from a Template

Use a CSV configuration file. Enter bare metal chassis information in the specified format and upload the file to import bare metal chassis.

Prerequisites

- Make sure 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.
- Parameters marked with * are required.
- You can enter a single IPMI address or IPMI range.

For IP ranges, use commas to separate addresses and use ^ to exclude a range. Example:

127.0.0.1-127.0.0.10,^127.0.0.2-127.0.0.3

- Bare metal reboot option:
 - Auto-reboot to obtain hardware information: Enter YES/Yes/yes/Y/y.
 - Manual reboot: Enter NO/No/no/N/n or leave blank.

Procedure

1. In the navigation pane, choose **Inventory > Bare Metal Management**.
2. Right-click the target bare metal cluster and choose **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, set the following parameters:
 - **Template**: Click **Download Template** and fill in bare metal chassis configuration according to the required format.
 - **Upload Template File**: Upload the file.
6. Click **OK**.

15.6.2 Bare Metal Chassis Status Management

Bare metal chassis have the following deployment status, status, and power status:

Deployment Status	Description
Unknown Hardware Info	The system failed to collect the hardware information of the bare metal chassis.

Deployment Status	Description
	Note: • Possible cause: You did not select the reboot bare metal chassis to obtain hardware information checkbox when adding the bare metal chassis. • Solution: Manually reboot the bare metal chassis or obtain hardware information.
PXE Booting	The deployment server (PXE server) is remotely instructing the bare metal chassis to boot from the PXE NIC and assigning a dynamic IP.
PXE Boot Failed	The bare metal chassis failed to boot from the PXE NIC. Note: Check these requirements: • Make sure no other DHCP services exist in the deployment network • Make sure PXE is enabled in BIOS for the NIC connected to the deployment network. For some models, make sure the PXE NIC is the first boot NIC or disable PXE on all NICs with higher boot priority. • Confirm the bare metal chassis boot mode is set to Legacy.
Available	The bare metal chassis is ready for creating bare metal instances.
Allocated	The bare metal chassis has been used to create a bare metal instance.
Rebooting	The bare metal chassis is rebooting.

Status	Power Status
Status includes: • Enabled • Disabled	Power status includes: • Power on • Rebooting • Power off

15.6.3 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 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.

15.6.4 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, click **Actions > Delete**.

Result



Note:

Deleting bare metal chassis will also delete bare metal instances created from these bare metal chassis. Proceed with caution.

15.7 Bare Metal Instance

15.7.1 New 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, click **Actions > New Bare Metal Instance**.
4. In the **New Bare Metal Instance** dialog, set the following parameters to complete the basic information:
 - **Name:** Set a name for the bare metal instance.
 - **Description:** Optional. Enter a description.
 - **Bare Metal Chassis:** Displays the bare metal chassis.
 - **Platform:** Only Linux is supported.
 - **Image:** Select an image for installing the OS.



Note:

- ISO format only.

- You need to upload the image to the standalone image storage in advance.
- Supports both custom platform OS and mainstream Linux distributions, including RHEL/CentOS series, Debian/Ubuntu series, and SUSE/openSUSE series.
- **Bare Metal Template:** Select a bare metal template to quickly generate pre-configured files for unattended batch OS installation. For more information, see [Bare Metal Template](#).
- **Auto Overwrite Data:** Choose whether to enable automatic data overwrite.

**Note:**

- When enabled, the original data will be automatically overwritten when bare metal instances are deployed.
- Otherwise, the installation of operating systems may pause. You need to access the console to manually configure disks.

5. Click **Next**.

6. Set the following parameters to complete system and network configurations:

You can configure system and network information for bare metal instances either individually or in batches.

- **Operating System:** Displays the selected OS.
- **Username:** Set a username. For batch configuration, default username is root.
- **Password:** Set a password.
- **Network Device:** Configure a business network for the bare metal instance.
 - **Device Type:** Choose NIC or NIC Bond.
 - **NIC Bond Name:** Set a name for the NIC bond, when selecting NIC bond as the device type.
 - **NIC Bond Type:** Select bond mode when selecting NIC bond as the device type. Options include Mode 1 (Active-Backup) and Mode 4 (LACP).
 - **NIC:** Select an NIC.
 - **Distributed Port Group:** Select a distributed port group.

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

15.7.2 Bare Metal Instance Status Management

Bare metal instances have the following state and deployment status:

State	Deployment Status
State includes: • Created • Powering on • Running • Stopped • Rebooting • Deleted • Error • Unknown	Deployment status includes: • Unprovisioned • Provisioning • Deployed

15.7.3 Delete a Bare Metal Instance

Prerequisites

(Optional) The platform provides deletion protection for bare metal instances. You can define how resources are deleted by customizing the deletion policy in system parameters. By default, the platform adopts a never delete policy for bare metal instances. Deleted resources are first moved to the recycle bin and retained until you manually expunge them. For more information, see [System Parameters](#).

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, click **Actions > Move to Recycle Bin**.

To delete multiple bare metal instances at once, go to the **Bare Metal Management** tab of the data center details page. Select the bare metal instances that you want to delete, then click **Bulk Action > Move to Recycle Bin**.

The delete button label changes based on the **Bare Metal Instance Deletion Policy**. When the deletion policy is set to immediate deletion, the button appears as "Delete". When the deletion policy is set to never delete, the button appears as "Move to Recycle Bin".

4. In the confirmation dialog, acknowledge the risk and click **OK**.

Result

Deleting a bare metal instance also powers off related bare metal chassis, which might affect your business continuity. Proceed with caution.

16 Storage Service

You can rapidly deploy a distributed storage cluster or seamlessly take over an external distributed storage system. You can centrally manage and utilize distributed storage within the platform, maximizing existing storage resources. The platform provides storage monitoring capabilities, allowing direct access to distributed storage overviews, node details, performance metrics, and disk health status.

16.1 Deploy Distributed Storage in 3 Steps

16.1.1 Step 1: Upload Installation Package

Prerequisites

- You have prepared the installation package of the distributed storage software.
- The platform management node must be added as a host.
- At least 2 distributed storage nodes have been added as hosts to the platform, and the storage network has been configured for storage nodes using Kernel adapters from hosts or distributed port groups. For more information about configuring storage network, see [Create a Kernel Adapter](#).
- To ensure the distributed storage can achieve MN HA together with the platform, we recommend setting up management node HA in MN Ops before deploying the distributed storage. For more information about setting up management node HA, see *Installation and Upgrade Tutorial*.

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, set the following parameters:

- **Server IP:** Display the IP address of the current management node.

In a dual management node environment, this field displays the IP address of the node where the VIP resides.

- **Storage Path:** Specify a directory path on the server to store uploaded installation packages and their extracted files.



Note:

- Ensure the directory has sufficient available space. Otherwise, installation may fail.
 - Do not use system directories such as `/`, `/dev`, `/proc`, `/sys`, `/usr/bin`, `/bin`, or `/opt`. Using system directories may cause server instability.
- **Upload By:** Select how to upload an installation package. Options include URL and Local File.

4. Click **OK**.

16.1.2 Step 2: Deploy Management Service

Procedure

1. On the **Storage Overview** page, click **Install**.
2. In the **Deploy Management Service** dialog, set the following parameters:
 - **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 single or dual management nodes. Enter the management node SSH port and password.
 - **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.
3. Review the configuration and click **OK**.

16.1.3 Step 3: Initialize Distributed Storage

Prerequisites

Before proceeding, you need to set up the storage network for distributed storage nodes using Kernel adapters from hosts or distributed port groups. For more information about configuring storage network, see [Create a Kernel Adapter](#).

Procedure

1. On the **Storage Overview** page, click **Initialize**.
2. In the **Initialize Distributed Storage** dialog, complete the initialization and storage configurations.
3. For **Initialization**, set the following parameters:
 - **Storage MN IP:** Select the management node IP of the distributed storage.
 - **Cluster:** Display the cluster location.

- **Admin Network:** Enter the admin network CIDR to manage and configure the storage cluster.
 - **Cluster Network:** Enter the cluster network CIDR to monitor data disks and synchronizing 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 to ensure time synchronization between all 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 to allow user services to 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 the manager (Mgr) and monitor (Mon) role to 3 nodes.

7. Specify a hostname for the distributed storage node.



Note:

- Length: 2 to 60 characters.
- Allowed characters: Uppercase and lowercase letters, numbers, and hyphens (-).
- No consecutive hyphens, nor start or end with a hyphen.
- When adding nodes, a suffix (-1, -2, -3, and so on) is automatically appended to ensure uniqueness.

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

What's next

Starting from NexaSphere 1.10.20, if you need to separately access the distributed storage management interface:

- 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

16.2 Take Over Existing Distributed Storage

16.2.1 Take Over Distributed Storage

Once the storage package installation, network configuration, and system initialization are complete on the server, or if the distributed storage is already running and in use, you can choose this way to take over the 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, set the following parameters:
 - **Storage MN IP**: Enter the management node IP of the distributed storage to be taken over.
 - **Storage Service Check**: Click **Check** to check if storage services exist and whether the connection can be established successfully.
 - **Username**: Enter the username for logging into the distributed storage platform.
 - **Password**: Enter the login password.
 - **Access Token**: Fill in this parameter as needed based on the detected storage service type.
4. Confirm the configuration and click **OK**.

What's next

Starting from NexaSphere 1.10.20, if you need to separately access the distributed storage management interface:

- 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

16.2.2 Cancel Takeover of Distributed Storage

Canceling the takeover of distributed storage will not affect any stored data, but you will no longer be able to 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.

16.3 Distributed Storage Resource Management

16.3.1 Storage Pool

16.3.1.1 Create a General Purpose Pool

On the main menu of NexaSphere, choose **Storage Service > Storage Pool**. On the **Storage Pool** page, click **Create Storage Pool**.

You can create three types of storage pools:

- Block Storage Pool
- Object Storage Pool
- File Storage Pool

Create a Block Storage Pool

Set the following parameters:

- **Name:** Set the name for the storage pool.

Naming rules: 1-1288 characters long. A name can contain Chinese characters, letters, digits, spaces, hyphens (-), underscores (_), periods (.), parenthesis (), colons (:), and plus signs (+).

- **Type:** Select **Block Storage**.
- **Role:** The default role is **Data Pool** and does not support modification.
- **Data Security Policy :**
 - **Type:** The default type is **Replicas** and does not support modification.
 - **Replicas:** Set the number of replicas for the storage pool in the 2-6 value range.



Note:

In production environments, we recommend setting at least 3 replicas to ensure data security.

- **Level:** Select the level of the failure domain (Server/Rack/Room) according to your topology plan.
- **Data Disk:** Select data disks based on the topology canvas.



Note:

- Selected data disks must meet the redundancy level requirements.
- Select data disks of similar sizes if possible.

Figure 16-1: Create a Block Storage Pool

< Create Storage Pool

Basic Info

Name *

Pool_1

Type *

Block Storage

Object Storage

File Storage

Role *

Data Pool

Data Security Policy

Type *

Replicas

Replicas *

☐ 2
☒ 3
☐ 4
☐ 5
☐ 6

Level *

☒ Server
☐ Rack
☐ Room

Data Disk *

osd.9 ×

osd.10 ×

osd.11 ×

Select Data Disk

Cancel

OK

Create an Object Storage Pool

Set the following parameters:

- **Name:** Set the name for the storage pool.

Naming rules: 1-128 characters long. A name can contain Chinese characters, letters, digits, spaces, hyphens (-), underscores (_), periods (.), parenthesis (), colons (:), and plus signs (+).

- **Type:** Select **Object Storage**.
- **Role:** Select the role of the storage pool (Data Pool, Index Pool, and Compound Pool).



Note:

1. Data Pool: Stores data.
2. Index Pool: Stores the index information of stored objects.
3. Compound Pool: Supports multi-purpose reuse and can be selected as an Index Pool or a Data Extra Pool in Storage Policy.

- **Data Security Policy :**
 - **Type:** Select data redundancy type (Replicas/EC).

- If you select Replicas, set the following parameters:
 - **Replicas:** Set the number of replicas for the storage pool in the 2-6 value range.

**Note:**

- Storage pools with Index Pool or Compound Pool role only support one redundancy policy, that is, Replicas.
- In production environments, we recommend setting at least 3 replicas to ensure data security.

- If you select EC, set the following parameters:
 - **EC Policy:** Set the EC policy for storage (Recommended/Custom).
 - **Recommended:** Select from six recommended values: 2+1, 4+2, 8+3, 4+2:1, 8+2:1, 16+2:1.
 - **Custom:** Customize the EC policy. Enter the number of data and parity blocks.

**Note:**

Positive integers only. Make sure that the number of data blocks is greater than the number of parity blocks, and parity blocks do not exceed 4.

**Note:**

- An EC policy consists of data blocks and parity blocks. Data blocks indicate the number of data shards, while parity blocks indicate the number of parity shards generated through the algorithm. Taking the 4+2 EC policy on the server level as an example. This policy ensures data availability even when 2 servers fail.
- Disk Utilization is displayed in real time. The formula for calculating disk utilization : data blocks/(data blocks + parity blocks).

- **Level:** Select the level of the failure domain (Server/Rack/Room) according to your topology plan.
- **Data Disk:** Select data disks to add based on the topology canvas.

**Note:**

- Selected data disks must meet the failure domain requirements of the data security policy.

- Select data disks of similar sizes if possible.

Figure 16-2: Create an Object Storage Pool

◀ Create Storage Pool

Name *

Type * Block Storage **Object Storage** File Storage

Role * ⓘ ☒ Data Pool ☐ Index Pool ☐ Compound Pool

Data Security Policy

Type * ☐ Replicas ☒ EC

EC Policy * ☒ Recommended ☐ custom
 ▼

Level * ☒ Server ☐ Rack ☐ Room

Data Disk * osd.9 × osd.10 × osd.11 × Select Data Disk

Cancel **OK**

Create a File Storage Pool

Set the following parameters:

- **Name:** Set the name for the storage pool.
 Naming rules: 1-1288 characters long. A name can contain Chinese characters, letters, digits, spaces, hyphens (-), underscores (_), periods (.), parenthesis (), colons (:), and plus signs (+).
- **Type:** Select **File Storage**.
- **Role:** Select the role of the storage pool (Data Pool/Metadata Pool).
- **Data Security Policy:**
 - **Type:** The **Replicas** type has been selected by default and you cannot modify it.
 - **Replicas:** Set the number of replicas for the storage pool in the 2-6 value range.



Note:

In production environments, we recommend setting at least 3 replicas to ensure data security.

- **Level:** Select the level of the failure domain (Server/Rack/Room) according to your topology plan.
- **Data Disk:** Select data disks to add based on the topology canvas.

**Note:**

- Selected data disks must meet the data redundancy level requirements.
- Select data disks of similar sizes if possible.
- To create a Metadata Pool, you need to use a raw SSD data disk.

Figure 16-3: Create a File Storage Pool

< Create Storage Pool

Basic Info

Name *

Pool_3

Type *

Block Storage

Object Storage

File Storage

Role *

☒ Data Pool
 ☐ Metadata Pool

Data Security Policy

Type *

Replicas

Replicas *

☐ 2
 ☒ 3
 ☐ 4
 ☐ 5
 ☐ 6

Level *

☒ Server
 ☐ Rack
 ☐ Room

Data Disk *

osd.9 ×

osd.10 ×

osd.11 ×

Select Data Disk

Cancel

OK

16.3.1.2 Manage a General Purpose Pool

On the main menu of NexaSphere, choose **Storage Service** > **Storage Pool**. Then, the **Storage Pool** page is displayed.

The following actions help to manage storage pools:

Action	Description
Create Storage Pool	Create storage pools.

Action	Description
Add Data Disk	Add one or more data disks to the storage pool.
Remove Data Disk	Remove a data disk to break the association with the storage pool. Hence a reduced storage pool capacity. Removing data disks may cause data losses. Proceed with caution. Note that you cannot remove data disks if: ◦ The storage pool is in a creating, deleting, or initializing status. ◦ The storage pool has a single replica. ◦ The data disk(s) to be removed has the only replica. ◦ When removing data disks in bulk, the rest of the data disks cannot meet the data security requirements. ◦ The cluster is enabled with Data Recovery and the pool capacity utilization may exceed the Backfill threshold after the removal. Note: If the pool capacity utilization is projected to exceed the Backfill threshold and you must remove the data disk(s), you can temporarily disable Data Recovery. The solution may cause data loss. Proceed with caution.
Set Recovery QoS	Choose the type of Recovery QoS for storage pools: Static QoS (Low Speed), Static QoS (Mid Speed), Static QoS (High Speed). When recovering the pool data, you can check data to recover, recovery rate, and remaining time on the General Purpose Pool page. • Low-Speed Recovery gives a higher priority to the business bandwidth. The recovery time is relatively long. Any hardware failures during recovery may reduce the data security level. We recommend that you choose Low-Speed Recovery in a production environment. • Mid-Speed Recovery gives the same priority to the business bandwidth and recovery bandwidth. The recovery time is medium. A saturated performance may increase the I/O latency. • High-Speed Recovery gives a higher priority to the recovery bandwidth. The recovery time is relatively short. A saturated performance may affect business performance. Note: You can specify whether to select high-speed recovery QoS for a storage pool in Global Setting. Default value: false. If set to true, you can select high-speed recovery when setting recovery QoS for a storage pool.

Action	Description
	Mid-Speed Recovery and High-Speed Recovery may affect your business and we do not recommend that you use them in a production environment. Contact technical support for a risk evaluation in advance if you have to choose Mid-Speed Recovery or High-Speed Recovery.
Check Data Consistency	Set check policy to execute data consistency checks in the storage pool. We support two check policies: - Default: The data consistency is checked once between 02:00 to 06:00 every day. - Per Day: The data consistency is checked once in a custom time range per day. Note: - A starting time equal to or greater than the ending time means the data consistency is checked between the starting time every day and the ending time on the next day.  We recommend that you avoid busy business hours to execute data consistency checks. - The check time accord with the server time. If the server time is not in sync with the browser time, data consistency checks will not be affected.
Rebalance Data Blocks	Manually rebalance the data blocks in the storage pool. Note:  You must enable Data Recovery for clusters and ensure that the storage pool is in a healthy state before performing this operation. - Rebalancing data blocks affects business performance. We recommend that you avoid busy business hours to perform this action.
Modify Data Security Policy	Modify the data security policy for storage pools. You are not provided with an option that does not meet the data security requirements. Replicas: For example, if the data redundancy level of the storage pool is server and the storage pool has 3 servers associated, you can set 2-3 replicas for the storage pool. You are not provided with an option that makes the pool capacity utilization exceed the Backfill threshold after the modification. If the pool capacity utilization has exceeded the Backfill threshold, you cannot set more replicas.

Action	Description
	- Erasure Code (EC): You can only switch from a collapsed EC policy to a standard EC policy. Note: - Modifying data security policy causes data migration which affects your business performance. We recommend that you avoid busy business hours to perform this action. - Note that the Backfill threshold is 85% by default. You can modify this value in Global Setting. - You can specify whether to set the number of replicas to 1 in  Global Setting. Default value: false. If set to true, you can set the number of replicas to 1 when creating a storage pool or modifying replicas. - If the storage pool operates in single-replica mode with insufficient capacity for replica modification, you cannot change the replica value. - When setting more replicas, you must verify the required total PG quantity after modification does not exceed the sum of maximum PG capacities across all data disks.
Delete	Delete the existing storage pools. Deleting a storage pool detaches all data disks from the pool. After deletion, the storage pool data cannot be recovered. Proceed with caution. Note: Delete the storage pools of Block Storage type: - Before deletion, you have to ensure that: - No block storage volumes exist in the storage pool. - The cluster where the storage pool is located is in the healthy  state. Delete the storage pool of Object Storage type: - For the storage pool chosen for initializing object storage, you have to check if it is associated with any resource, such as Storage Policy, Object Gateway, Object User, and Bucket. If the answer is no, you can delete it. Deleting the storage pool will cause the loss of the object storage system resources pool. After deletion, you

Action	Description
	cannot get the object storage service and access the data within. Proceed with caution. - For the storage pool chosen for initializing object storage, you have to check if it is associated with any resource, such as Storage Policy, Object Gateway, Object User, and Bucket. If the answer is yes, you cannot delete it. - For the storage pool not chosen for initializing object storage, you have to check if it is associated with any Storage Policy. If the answer is yes, you cannot delete the pool. Delete the storage pool of File Storage type: - If a storage pool is associated with a file system, it cannot be deleted. - If a storage pool is in a creating, scaling-in, scaling-out, updating, or deleting status, it cannot be deleted.

16.3.2 Storage Node

16.3.2.1 Add a General Purpose Storage Server

On the main menu of NexaSphere, choose **Storage Service > Storage Node**. On the **Server** page, click **Add Server**.

Adding a general purpose storage server involves five steps. Set the following parameters to complete the server configurations.

Step One: Basic Configurations

- **Server IP Address:** Enter the server IP address. You can specify either a single IP address or IP address range. We support adding multiple servers in bulk.
- **Type:** Select server type. Two server types are supported:
 - Storage Server:
 - Provides storage pools with hard disks that can be used as data disks.
 - Supports five roles: Management, Monitor, Block Storage Gateway, Object Storage Gateway, and File Storage Gateway.
 - Storage Gateway Server:

- Hosts various interfaces and clients. The system only manages the server gateways and does not manage hard disks on the server.
- Supports only one role: Block Storage Gateway.
- **Role:** Configure the role of servers. Five roles are supported:
 - Admin Role (Management):
 - Responsible for the collection and management of the runtime status of the cluster and manages the distributed storage cluster as the management node in multiple ways, such as GUI and API.
 - We recommend that you deploy at least 2 admin roles to meet the high availability requirement.
 - Monitor Role:
 - Responsible for monitoring the cluster storage data and maintaining overall status of the cluster, including metadata such as data mapping and cluster authentication.
 - We recommend that you deploy an odd number of monitor roles ($3+2*N$, $N \geq 0$) to meet the high availability requirement.
 - Block Storage Gateway:
 - Responsible for the access between the server and the storage cluster through Block interface.
 - By default, this role is selected for a storage server.
 - By default, this role is selected for a gateway server. And a storage gateway server only supports this role.
 - Object Storage Gateway:
 - Responsible for the access between the server and the storage cluster through Object interface.
 - To use object storage service, you need to select this role.
 - On an object storage gateway server, you can turn on the object gateway to provide the S3 protocol and gateway services.
 - File Storage Gateway:
 - To use file storage service, you need to select this role.
 - On a file storage gateway server, you can create a file gateway to provide file storage access protocols such as SMB and NFS.

**Note:**

- When you add a server for the first time, three roles including Management, Monitor, and Block Storage Gateway, are selected by default. The Block Storage Gateway role can be deselected, while Management and Monitor roles are required.
- For subsequent server additions, you can add storage servers without roles.
- Deploy at least three storage servers with Management, Monitor, and Block Storage Gateway roles in a cluster.

Figure 16-4: Basic configurations

< Add Server

Basic Configurations ○

Environment Configurations ●

Network Configurations ●

Confirm ●

Install ●

Server IP Address *

☒ IP Address ☐ IP Address Range



[+ Add IP Address](#)

Type * ⓘ

☒ Storage Server ☐ Storage Gateway Server

Role * ⓘ

☒ Management ☒ Monitor ☒ Block Storage Gateway ☐ Object Storage Gateway ☐ File Storage Gateway

Cancel [Next >](#)

Step Two: Environment Configurations

- **SSH Username:** Enter the SSH username for the server. Default: **root**.
- **SSH Password:** Enter the SSH password. The system uses this password only for password-free login configurations and does not store the password.
- **Port:** Enter the server port number. Default port: **22**.
- **Server Name:** (Optional) Specify a server name.

Naming rules: 1-63 characters long. The name can contain lower-case letters (a-z), digits (0-9), periods (.), and hyphens (-). Avoid starting with a hyphen or number as well as ending with a hyphen.

**Note:**

- If you do not set a server name, ensure the server name-to-IP mapping is preconfigured in `/etc/hosts` file. The system uses the existing server name after server addition.
 - If you set a new server name, it overwrites the existing name-to-IP mapping in `/etc/hosts` file.
 - When you add servers in bulk, the names of these servers will end with a suffix, that is, the last part of their IP address (0-254), to distinguish these servers, for example, server-24.
- **Time Sync Service:** Choose to enable or disable the time synchronization service.
If you enable this setting, the system synchronizes the newly-added server's clock with other servers in the cluster.
 - **Password-Free Login:** If you enable this setting, the system configures password-free logins to the server with the SSH username and password.

Figure 16-5: Environment configurations

< Add Server

Basic Configurations

Environment Configurations

Network Configurations

Confirm

Install

SSH Username *

SSH Password *

Port *

Server Name

Time Sync Service

Password-Free Login

root

22

172.26.51.12

☒ Set Server Name ☐ Not Set Server Name

☒

☒

The system uses the password only for password-free login configurations and does not store the password.

Turn on the button to ensure the time consistency of the newly-added servers with other servers in the cluster.

Turning on the button means that you allow the system to configure password-free logins to the server with the SSH username and password.

Cancel < Previous **Next >**

Step Three: Network Configurations

- **Admin IP:** Sets the IP address the management network which manages and configures storage clusters. The default admin IP is the server IP address.

- **Public IP:** Sets the IP address of the public cluster network which facilitates interaction between block storage gateways and storage pools.
- **Cluster IP:** Sets the IP address of the cluster internal network which monitors data disks across cluster servers and synchronizes replicas.

**Note:**

Skip setting Cluster IP when you add a storage gateway server.

Figure 16-6: Network configurations

< Add Server

Basic Configurations

Environment Configurations

Network Configurations

Confirm

Install

Network Configurations

	Admin IP	Public IP *	Cluster IP *
1	172.26.51.12	172.26.51.12	172.26.51.12

Cancel < Previous Next >

Step Four: Confirm

Review the information of the server to add. You can navigate back to modify configuration details if needed.

Figure 16-7: Confirm

< Add Server

Basic Configurations

Environment Configurations

Network Configurations

Confirm

Install

Basic Configurations

Server IP Address : 172.26.51.12

Type : Storage Server

Role : Management, Monitor, Block Storage Gateway

Environment Configurations

SSH Username : root

SSH Password : *****

Port : 22

Sever Name : 172.26.51.12

Time Sync Servi... : Enable

Password-Free ... : Enable

Network Configurations

	Admin IP	Public IP	Cluster IP
1	172.26.51.12	172.26.51.12	172.26.51.12


Cancel < Previous Next >

Step Five: Install

Check the server deployment progress. If you exit this page during installation, the process continues in the background. You can check the process via the operation log at any time.

**Note:**

- You cannot exit the page when adding a server for the first time.
- Avoid refreshing the browser when adding a server for the first time.

Figure 16-8: Install

16.3.2.2 Manage a General Purpose Storage Server

On the main menu of NexaSphere, choose **Storage Service > Storage Node > General Purpose Node**. Then, the **General Purpose Node** page is displayed.

The following actions help to manage general purpose storage servers.

Action	Description
Add Server	Add one or more servers. Note: If a monitoring server in the disconnected state exists in the cluster, you cannot add a new server.
Delete Server	Deleting a server stops the services provided by the server and deletes all relevant data. Deleted data cannot be recovered. Proceed with caution. Note: To delete a server requires meeting the following requirements: - The server is in the connected state. - The cluster where the server is located is in the healthy status. - No data disks are running in the server. - The server is not added as an Object Storage Gateway.

Action	Description
	The server is not added as a Block Storage Gateway.

16.3.3 Data Disk

16.3.3.1 Create a Data Disk on General Purpose Node

On the main menu of NexaSphere, choose **Storage Service** > **Data Disk**. On the **Data Disk** page, click **Create Data Disk**.

Set the following parameters:

- **Select Hard Disk:** Select the hard disk to add as a data disk.



Note:

You can add hard disks that are healthy, free, and of at least 25 GB size as data disks.

- **Cache Partition:** Choose whether to set cache partition for data disks.
 - **Auto Mode:** If you select auto mode, the system automatically attaches cache partitions provided by servers to the selected hard disks.
 - **Manual Mode:** If you select manual mode, you can manually select cache partitions for each hard disk that has been selected.



Note:

If available cache partitions are insufficient for the selected hard disks when you enable this parameter, part of the hard disks fail to be added as data disks.

Figure 16-9: Create a data disk

< Create Data Disk

Select Hard Disk * ⓘ

/dev/sde ×

Select a hard disk.

Cache Partition ⓘ

☒ ☐

☐ Auto Mode
 ☒ Manual Mode

Partition Setting	Hard Disk	Server	Cache Partition
	/dev/sde	172.26.50.182	⚠ No available cache partition. Check

16.3.3.2 Manage a Data Disk on General Purpose Node

On the main menu of NexaSphere, choose **Storage Service > Data Disk**. Then, the **Data Disk** page is displayed.

The following actions help to manage data disks.

Action	Description
Create Data Disk	Create one or more data disks.
Set Maintenance Mode	Enable or disable the maintenance mode for data disks. A data disk in the maintenance mode is not involved in data block rebalances. Note: - The maintenance mode stops services provided by and access to the data disk. To ensure system availability, we do not recommend that you put a data disk in the maintenance mode for a long time.  A maintenance mode does not stop data migrations on the data disk. To stop data migrations, manually disable Data Recovery in Global Setting. - Disable the maintenance mode and enable Data Recovery manually after you finish the disk maintenance. The services and data access automatically recover after you turn off the maintenance mode.
Delete	Deleting a data disk stops services provided by the data disks and deletes all relevant data. The deleted data cannot be recovered. Proceed with caution. Note:  To delete a data disk requires meeting these requirements: - The associated cache disk must be in a healthy state. - The data disk must not be part of any storage pool.

16.3.4 Physical Hard Disk

16.3.4.1 Scan Hard Disks on General Purpose Node

On the main menu of NexaSphere, choose **Storage Service > Physical Hard Disk**. On the **Hard Disk** page, click **Scan** and all hard disks on storage servers and their information will be displayed in the list.

Figure 16-10: Scan hard disks

Hard Disk
A hard disk is the physical unit of a data disk. All hard disks on storage servers are scanned and displayed in the list. Healthy free disks can be added as data disks.

General Purpose Node | High-Performance Node

Scan Bulk Action Search

Identifier	State	Use	Medium	Total Capacity	Server	Drive Path	Disk Light	Actions
wwn-0x000f47c73af9f05f	Healthy	Data Disk	HDD	150 GB	172.26.51.216	/dev/sdc	-	...
virtio-Sacard792fb8b4ab4884f	Healthy	System Disk	HDD	300 GB	172.26.51.216	/dev/vda	-	...
wwn-0x000fc17db8f4a54e	Healthy	Data Disk	HDD	150 GB	172.26.51.216	/dev/sdb	-	...
wwn-0x000f0998d9a8ab35	Healthy	Data Disk	HDD	150 GB	172.26.51.216	/dev/sda	-	...
wwn-0x000f990bb205b247	Healthy	Data Disk	HDD	150 GB	172.26.50.55	/dev/sdb	-	...
wwn-0x000fa8623c25ae32	Healthy	Data Disk	HDD	150 GB	172.26.50.55	/dev/sdc	-	...
virtio-a0244d30121844e8ab1c	Healthy	System Disk	HDD	300 GB	172.26.50.55	/dev/vda	-	...
wwn-0x000fc0e12d554431	Healthy	Data Disk	HDD	150 GB	172.26.50.55	/dev/sda	-	...
wwn-0x000fe3d4cda1de29	Healthy	Data Disk	HDD	150 GB	172.26.50.182	/dev/sdb	-	...
wwn-0x000f731f9f40a04	Healthy	Data Disk	HDD	150 GB	172.26.50.182	/dev/sda	-	...

Item 1 to 10. Total: 13. < 1 2 > 10 Items/Page

**Note:**

- Healthy free disks can be used for cache partitioning (SSD recommended) or added as data disks. If you set cache partition, ensure that each partition has a minimum capacity of 50 GB.
- A scanned **unknown disk** implies that the disk contains unrecognized partitions.
 - For NexaSphere 4.2.0 and earlier versions, run the `wipefs -af /dev/sdX` command to manually clean up partitions before rescanning the disk.
 - For NexaSphere 4.2.0 and later versions, go to the **Hard Disk** page and click **Initialize Hard Disk** to clean up partitions directly.
- In some hardware environments, newly-added hard disks may not be detected. Reboot the server or contact official technical support for assistance.

16.3.4.2 Manage Hard Disks on General Purpose Node

On the main menu of NexaSphere, choose **Storage Service** > **Physical Hard Disk**. Then, the **Hard Disk** page is displayed.

The following actions help to manage hard disks.

Action	Description
Scan	Scan and list all hard disks in the server and their use.
Set Cache Partition	You can set cache partitions for one or more healthy free disks.

Action	Description
	Note: • Each partition has a minimum capacity of 50 GB.。 • You can set up to 36 partitions. • If you fail to perform bulk action, select free disks of the same total capacity and try to set cache partition again.。
Clean up Cache	We support cleaning up cache for cache disks in the healthy state. Note: • Cleaning up cache means cleaning up all partitions in the cache disk. The clean-up cannot be recovered. Proceed with caution. • You can reset cache partition or add hard disks as data disks. • If the partitions of a cache disk is currently in use by data disks, you cannot clean up the cache disk. You need to delete the associated data disks first before trying to clean up cache again.
Initialize Hard Disk	Initialize one or more disks in the healthy state whose use is unknown. Note: • After initialization, the use of a hard disk will change from Unknown Disk to Free Disk. • Initializing hard disks will erase all existing partitions. This operation may cause permanent data loss. Proceed with caution.
Disk Light	Enable or disable the disk light to quickly locate the hard disk. Note: • You cannot light up a system disk, virtual disk, or offline disk. • We recommend you use the hard disk compatible with our platform , for example, HGST HUS728T8TALE6L4 or ST2000DM001-1ER164 and so on.

18 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.

Network 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 is used to collect logs of the management node. You can add a log server to the platform and use the collected logs to locate errors and exceptions. This makes your O&M more efficient.

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.