NEXAVM CLOUD

TECHNICAL WHITE PAPER

Contents

1 Overview

The mobile cloud era is changing line-of-business (LOB) expectations of IT. For IT organizations to securely deliver the anticipated improvements in service quality and speed, a Software-Defined Data Center (SDDC) approach is required. Driven by the breakthrough of the 'software-defined' architecture. Virtual machine (VM), virtual network and corresponding storage could be provisioned and reconfigured in a high-speed and automated way, without any compromise with non-dynamic hardware architecture. Software-Defined Data Center (SDDC) enables users to focus on application, while required IT architecture resource scheduling is accomplished accordingly by the software dynamically, that is, the hardware resource scheduling is accomplished in a software-based method.

NexaVM Cloud is the next-generation software defigned solution mainly for future-oriented, smart data centers. Also, it manages multiple compute, storage, and network resources in data centers by providing flexible and comprehensive APIs. You can quickly set up your own smart cloud data center by using NexaVM Cloud, or construct flexible cloud usage scenarios, such as Cloud Desktop, PaaS, and SaaS, on NexaVM Cloud.

Figure 1-1: NexaVM Cloud Framework

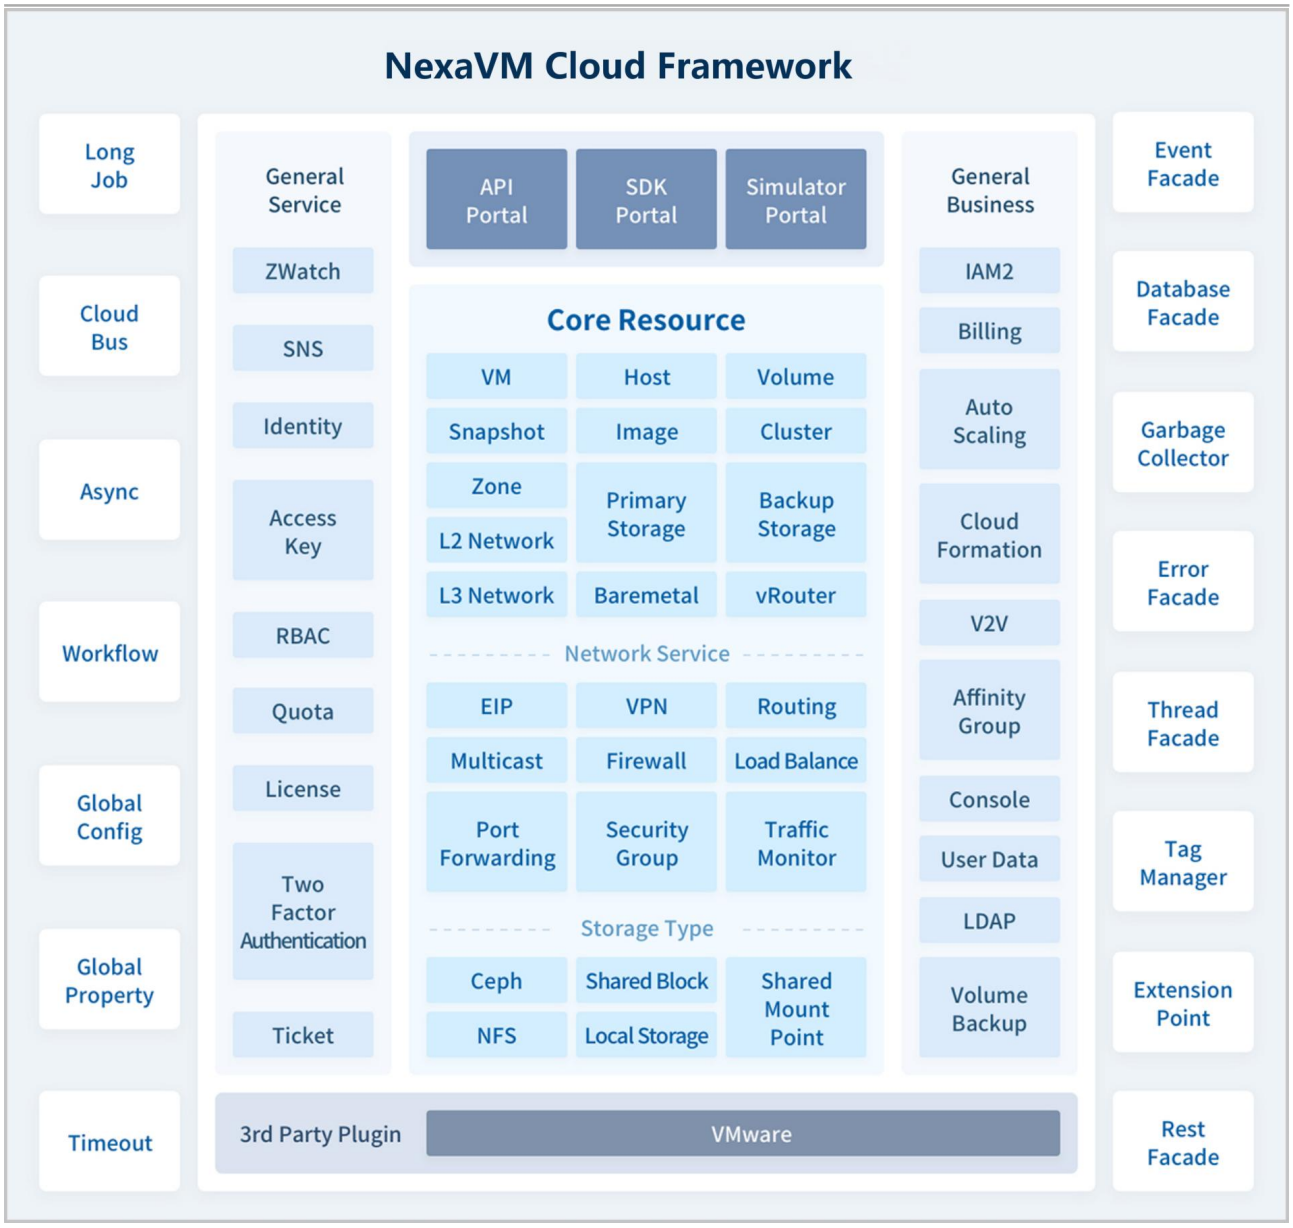
NexaVM Cloud Framework
Long Job
Cloud Bus
Async
Workflow
Global Config
Global Property
Timeout
General Service
ZWatch
SNS
Identity
Access Key
RBAC
Quota
License
Two Factor Authentication
Ticket
API Portal
SDK Portal
Simulator Portal
Core Resource
VM
Host
Volume
Snapshot
Image
Cluster
Zone
Primary Storage
Backup Storage
L2 Network
L3 Network
Baremetal
vRouter
Network Service
EIP
VPN
Routing
Multicast
Firewall
Load Balance
Port Forwarding
Security Group
Traffic Monitor
Storage Type
Ceph
Shared Block
Shared Mount Point
NFS
Local Storage
General Business
IAM2
Billing
Auto Scaling
Cloud Formation
V2V
Affinity Group
Console
User Data
LDAP
Volume Backup
3rd Party Plugin
VMware
Event Facade
Database Facade
Garbage Collector
Error Facade
Thread Facade
Tag Manager
Extension Point
Rest Facade

2 Product Profiles

2.1 NexaVM Cloud Functional Architecture

The functional architecture of NexaVM Cloud is shown in *Figure 2-1: NexaVM Cloud Functional Architecture*.

Figure 2-1: NexaVM Cloud Functional Architecture

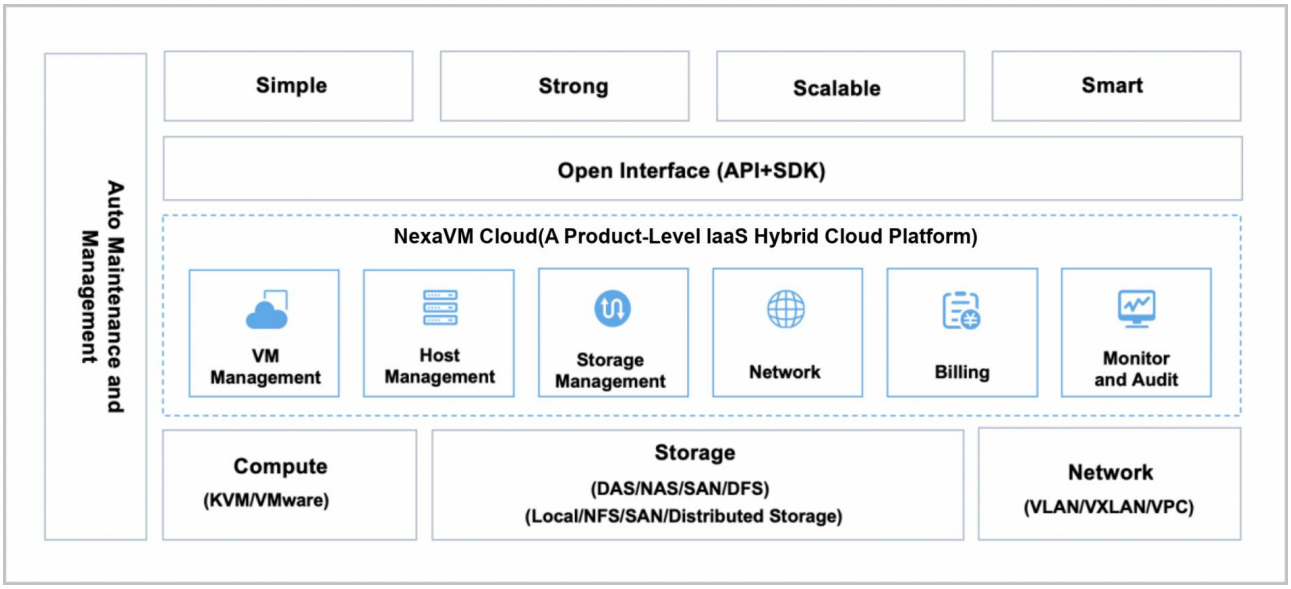

NexaVM Cloud helps enterprises better manage their infrastructure resources, such as the compute, storage, and network resources, in their data centers. The bottom layer of NexaVM Cloud supports both KVM and VMware virtualization technologies. In addition, NexaVM Cloud supports various storage types, such as DAS, NAS, SAN, and DFS. To be more specific, local storage,NFS storage, SAN storage, and distributed block storage are supported. NexaVM Cloud also supports various network models, such as VLAN and VXLAN.

NexaVM Cloud communicates with the MariaDB database and other service modules by using a message bus, providing diversified features such as VM instance management, host management, storage management, network management, billing management, and real-time monitoring. In addition, NexaVM Cloud provides Java SDKs and Python SDKs, and allows you to schedule and manage resources by using RESTful APIs. With NexaVM Cloud, you can build a private cloud that is simple, strong, scalable, and smart.

Highlights of NexaVM Cloud functional architecture:

1. Asynchronous Architecture: asynchronous message, asynchronous method, and asynchronous HTTP call

- NexaVM Cloud connects various services by using a message bus. When a service calls another service, the source service sends a message to the destination service, registers a callback function, and then returns back immediately. Once the destination service finishes the task, it gives a feedback on the task result by triggering the callback function that was registered by the source service. Asynchronous messages can be processed in parallel.
- Services in NexaVM Cloud communicate with each other through asynchronous messages. The components and plugins inside each service are also called by using the asynchronous method, which is the same as calling asynchronous messages.
- Every plugin in NexaVM Cloud has a corresponding agent. NexaVM Cloud puts a callback URL in the HTTP header of every request. Therefore, agents can send responses to the URL of the caller when tasks are finished.
- Based on the asynchronous message, asynchronous method, and asynchronous HTTP call, NexaVM Cloud builds a layered architecture to ensure that asynchronous operations can be performed on all components.
- Based on the asynchronous architecture, NexaVM Cloud with a single management node can process tens of thousands of concurrent API requests per second, and simultaneously manage tens of thousands of servers and hundreds of thousands of VM instances.

2. Stateless Service: A single request does not rely on other requests.

- In NexaVM Cloud, requests sent by compute node agents, storage agents, network services, console proxy services, and configuration services can be processed without relying on other requests. The sent requests contain all the required information, and related nodes do not need to maintain or store any information.
- NexaVM Cloud authenticates resources such as management nodes and compute nodes through consistent hashing ring by using their UUIDs as the unique ID. Because of the consistent hashing ring, a message sender does not need to know which service instance is about to handle the message. Services do not need to maintain or exchange information about what resources they are managing. All the services need to do is to handle the incoming messages.
- Little information is shared among NexaVM Cloud management nodes. Therefore, a minimum of two management nodes can meet the requirements of high availability and scalability.
- The stateless service mechanism makes the system more robust. Restarting the server will not lose any state information. This also simplifies the scaling out and scaling in of a data center.

3. Lock-free Architecture: consistent hashing algorithm

- The consistent hashing algorithm guarantees all messages of the same resource are always handled by the same service instance. In this way, messages are congregated to a specified node, reducing the complexity of synchronization and concurrency.
- NexaVM Cloud uses work queue to avoid lock contention. Serial tasks are stored in memory as work queues. Work queues can process any operation of any resource in parallel to improve system concurrency.
- The queue-based lock-free architecture enables tasks to run in parallel, thereby improving the system performance.

4. In-Process Microservices Architecture: microservices decoupling

- NexaVM Cloud uses a message bus to isolate and control various services, such as VM instance services, identity authentication services, snapshot services, volume services,network services, and storage services. All microservices are enclosed in the same process of a management node. These services communicate with each other through the message bus. After all messages are sent to the message bus, the destination service is selected by the consistent hashing ring for message forwarding.
- In-process microservices provide a star-like architecture, ensuring every service in microservices to run independently. This architecture also decouples the highly centralize d control business, and achieves a high degree of autonomy and isolation of the system. Failure of any service does not affect other components. This effectively guarantees the system reliability and stability.

5. Versatile Plugin System: supports horizontal expansion of plugins

- In NexaVM Cloud, every plugin provides services independently. Any newly added plugin has no impact on other existing plugins.
- NexaVM Cloud concludes plugins into two patterns: strategy pattern and observer pattern. Strategy pattern plugins will inherit parent-class interfaces and then perform specific implementations. Observer pattern plugins will register a listener to monitor event changes of the internal business logic in an application. Once an event is detected inside the application, the observer pattern plugins will respond to this event automatically and execute a piece of code to affect the corresponding business flow.
- NexaVM Cloud supports horizontal expansion of plugins. The Cloud can be quickly upgraded, and the overall system architecture still remains robust.

6. Workflow Engine: sequence-based management, rollback on errors

- NexaVM Cloud clearly defines every workflow by using XML files. Every flow can be rolled back on errors. A workflow can roll back all prior executed steps and clean up the garbage resources during the execution when an error happens in a step.
- Each workflow can contain a sub-flow to decouple the business logic further.

7. Tag System: extends the business logic and adds resource properties

- NexaVM Cloud uses system tags and plugins to extend the original business logic.
- You can use tags to group your resources and search for resources with specific tags.

8. Cascade Framework: supports cascading operations on resources

- NexaVM Cloud uses a cascade framework to perform cascading operations on resources. The cascade framework allows an operation to be cascaded from one resource to other resources. For example, the operation of uninstalling or deleting a resource can be cascaded to the descendant resources.
- Resources can join in a cascade framework through a plugin. Joining or quitting the cascade framework will not affect other resources.
- The cascading mechanism makes the configuration of NexaVM Cloud more flexible and simple, meeting the requirements of resource configuration changes.

9. Full Automation By Ansible: automated deployment by agentless Ansible

- Being seamlessly integrated with Ansible (which is agentless), NexaVM Cloud can automatically install dependencies, configure physical resources, and deploy agents. This whole process is transparent to users and requires no additional intervention. You can upgrade your agents simply by reconnecting the agents.

10.Comprehensive Query API: Every property of every resource can be queried.

- NexaVM Cloud supports millions of query conditions, comprehensive query APIs, and any way of condition combinations.

2.2 NexaVM Cloud Resource Model

NexaVM Cloud is essentially a configuration management system for resources in the Cloud. The following figure describes the resource model managed by NexaVM Cloud, as shown in *Figure 2-2: NexaVM Cloud Resource Model*.

Figure 2-2: NexaVM Cloud Resource Model

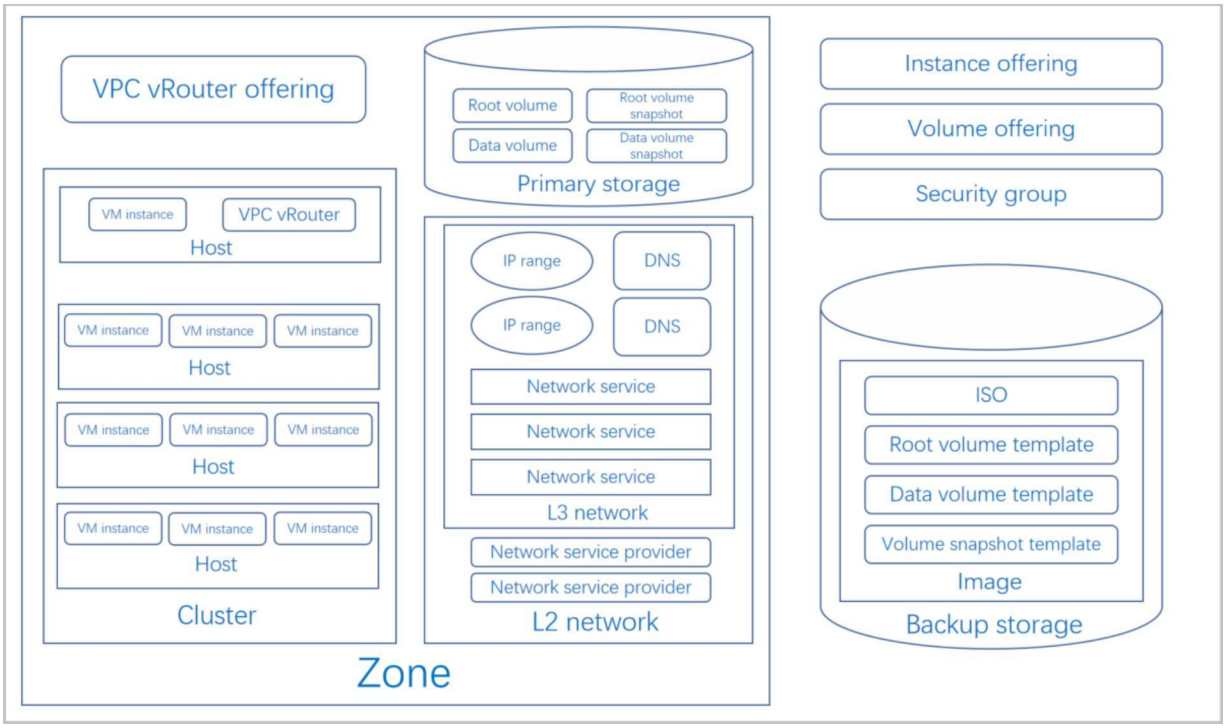

NexaVM Cloud mainly has the following resources:

- Zone: the largest resource scope defined in NexaVM Cloud. A zone is a logical group of resources, such as clusters, L2 networks, and primary storages.
- Cluster: a logical group of analogy hosts (compute nodes).
- Host: also known as a compute node, is a physical server that provides VM instances with compute, network, and storage resources.
- Primary storage: a storage system that stores VM disk files, including root volumes, data volumes, root volume snapshots, data volume snapshots, and image caches. Supported primary storage types include local storage, NFS, SharedMountPoint, SharedBlock, and Ceph.
- Backup storage: a storage system that stores image templates. Supported backup storage types include ImageStore, SFTP, and Ceph.
- VXLAN pool: an underlay network in VXLAN. You can create multiple VXLAN overlay networks (VXLAN) in a VXLAN pool. The overlay networks can operate on the same underlay network device. Supported VXLAN pool types include software SDN and hardware SDN.
- L2 network: a layer 2 broadcast domain used for layer 2 isolation. Generally, L2 networks are identified by names of devices on the physical network. Supported L2 network types include L2NoVlanNetwork, L2VlanNetwork, VxlanNetwork, and HardwareVxlanNetwork.

- L3 network: a collection of network configurations for VM instances, including the IP range, gateway, DNS, and network services.
- Instance offering: a specification that defines the number or size of CPU, memory, disk bandwidth, and network bandwidth for a VM instance.
- Disk offering: a specification of a volume, which defines the size of a volume and how the volume will be created.
- VM instance: a virtual machine instance running on a host. A VM instance has its own IP address to access public network and run application services. VM instances are core components of NexaVM Cloud.
- Image: an image template used by a VM instance or volume. Image template includes root volume images and data volume images. The types of root volume image include ISO and Image, while the type of data volume image is Image.
- Root volume: the system disk where the VM operating system is installed.
- Data volume: the data disk that provides additional storage for a VM instance.
- Snapshot: a point-in-time capture of data in a disk. Snapshots are captured incrementally.
- Network service module: a module for providing network services. This resource is hidden in the UI.
- Network service: provides various network services for VM instances, including VPC firewall, security group, virtual IP (VIP), elastic IP (EIP), port forwarding, load balancing, IPsec tunnel, and flow monitoring.
- VPC firewall: manages north-south traffic of the VPC network. You can manage the network access policy by configuring rule sets and rules.
- Security group: provides L3 network firewall control over the VM instances, and controls TCP, UDP, and ICMP data packets for effective filtering. You can use a security group to effectively control specified VM instances on specified networks according to specified security rules.
- vRouter offering: an instance offering that defines the number of vCPU cores, memory size, image, management network, and public network configuration settings of VPC vRouters.
- VPC vRouter: a router created directly from a vRouter offering. VPC vRouter, which has a public network and a management network, is the core of VPC. VPC vRouters provide various network services, including DHCP, DNS, SNAT, route table, EIP, port forwarding, load balancing, IPsec tunnel, dynamic routing, multicast routing, VPC firewall, and Netflow.

The resource relationships in NexaVM Cloud are as follows:

- Parent-child: A resource can be the parent or child of another resource. For example, a host is the child resource of a cluster and the parent resource of a VM instance.
- Sibling: Resources sharing the same parent resource are siblings. For example, clusters and L2 networks are sibling resources because all of them are child resources of a zone.
- Ancestor-descendant: A resource can be the lineal ancestor or lineal descendant of another resource. For example, a cluster is the ancestor resource of a VM instance, while a host is a descendant resource of a zone.
- Friend: Resources that do not have the above three relationships but still need to cooperate with each other in some scenarios are friends. For example, primary storages and backup storages are friends. Also, zones and backup storages are friends.

Note:

Relationship between primary storages and backup storages:

- When you create a VM instance, the primary storage needs to download images of the VM instance as caches from the backup storage.
- When you create an image, the primary storage needs to copy the root volume to the backup storage and save it as a template.

The following properties are common to almost all resources in NexaVM Cloud:

- UUID: the universally unique identifier. NexaVM Cloud uses version 4 UUIDs to uniquely identify

  a resource.
- Name: a human readable string that is used to identify resources. Names can be duplicated and are usually required.
- Description: also known as a brief introduction that is used to briefly describe a resource. Description is usually optional.
- Creation date: the date and time when a resource was created.
- Last operation date: the date and time when a resource was updated last time.

Resources in NexaVM Cloud support full or partial Create, Read, Update, Delete (CRUD) operations.

- Create: create or add a new resource.
- Read: read or query information about a resource.
- Update: update information about a resource.

- Delete: delete a resource. Due to the cascade framework provided by NexaVM Cloud, if a parent resource is deleted, its associated child resources and descendant resources will also be

  deleted.

2.2.1 Resource Center

2.2.1.1 Resource Pool

2.2.1.1.1 VM Instance

A VM instance is a virtual machine instance running on a host. A VM instance has its own IP address and can access public networks and run application services.

2.2.1.1.2 Volume

A volume provides storage space for a VM instance. Volumes are categorized into root volumes and data volumes.

Note:

Volume management mainly involves data volumes.

Concepts

- Root volume: A root volume provides support for the system operations of a VM instance.
- Data volume: A data volume provides extended storage space for a VM instance.

Considerations

- A volume attached to a VM instance cannot be attached to a VM instance of another hypervisor type. For example, a KVM VM volume cannot be attached to a VMware VM instance.
- The storage space that a volume takes up is calculated based on the virtual size of the volume. When a volume is created, the virtual size of the volume is deducted. The actual storage space that the volume takes up is small upon creation while increases along with growing data writes.
- A non-shared volume can be attached to only one VM instance. A shared volume can be created on a Ceph primary storage and SharedBlock primary storage and allows simultaneous access from multiple VM instances.
- A root volume is an integral part of a VM instance. It cannot be detached from a VM instance.
- A data volume attached to a VM instance can be detached from the VM instance and then attached to another VM instance of the same hypervisor type.

- When multiple primary storage are available, you can specify a primary storage to create a volume. If you do not specify a primary storage when you create a volume, note that:
  - For LocalStorage primary storage, the volume is created on the primary storage with the largest available space by default.
  - For NFS primary storage, the volume is created on a random primary storage by default.
  - For the combination of LocalStorage and NFS primary storage or the combination of LocalStorage and SharedMountPoint (SMP) primary storage, by default, the volume is created on a primary storage that does not store the root volume of the VM instance.
- You can set the QoS for a data volume to limit its disk bandwidth. Note that an excessively low QoS may greatly lower I/O performance.

2.2.1.1.3 Image

An image is a template file used to create a VM instance or volume. Images are categorized into system images and volume images.

- A system image contains the operating system your business runs on. You can use it to create VM instances. A system image can be of the ISO or Image type.
- A volume image contains only your business data. You can use it to create volumes. A volume image can only be of the Image type.
- Image-type images can be in the raw, qcow2, or vmdk format.
- Images are stored on backup storage. When an image is used to create a VM instance or volume, the image is downloaded to a primary storage and cached there.

2.2.1.1.4 Instance Offering

An instance offering defines the number of vCPU cores, memory size, network bandwidth, and other configuration settings of VM instances.

2.2.1.1.5 Disk Offering

A disk offering defines the capacity and other configuration settings of volumes.

A disk offering can be used to create root volumes and data volumes.

2.2.1.1.6 GPU Specification

A GPU specification defines the frame per second (FPS), video memory, resolution, and other configuration settings of a physical or virtual GPU. GPU specifications are categorized into physical GPU specifications and virtual GPU specifications. NexaVM Cloud supports GPUs installed

on hosts. In NexaVM Cloud, you can directly assign a single physical GPU to a VM instance, or share it among multiple VM instances running on the same host.

2.2.1.1.7 Auto-Scaling Group

NexaVM Cloud offers auto-scaling capabilities that can automatically add or remove VM instances from an auto-scaling group (ASG) in response to load balancing of VM instances, your business load changes, and predefined scaling policies. With the auto scaling service, you can better

leverage the Cloud resources, reduce the O&M costs, and ensure smooth business operations. Currently, the auto scaling service is applicable to KVM VM instances.

Scaling Mode

Currently, the Cloud supports the following two types of scaling mode:

1. **Auto Scaling**

   - Auto scaling includes elastic scale-out and elastic scale-in. For elastic scale-out, when your businesses are growing, VM instances will be automatically added to ensure your business continuity. For elastic scale-in, if your businesses decrease, VM instances will be automatically reduced.
   - The Cloud Monitoring feature allows you to trigger the auto scaling mode and specify an endpoint to receive alarms. Supported endpoints include system endpoint, email, DingTalk, HTTP application, short message service (SMS), and Microsoft Teams.

Figure 2-3: Auto Scaling

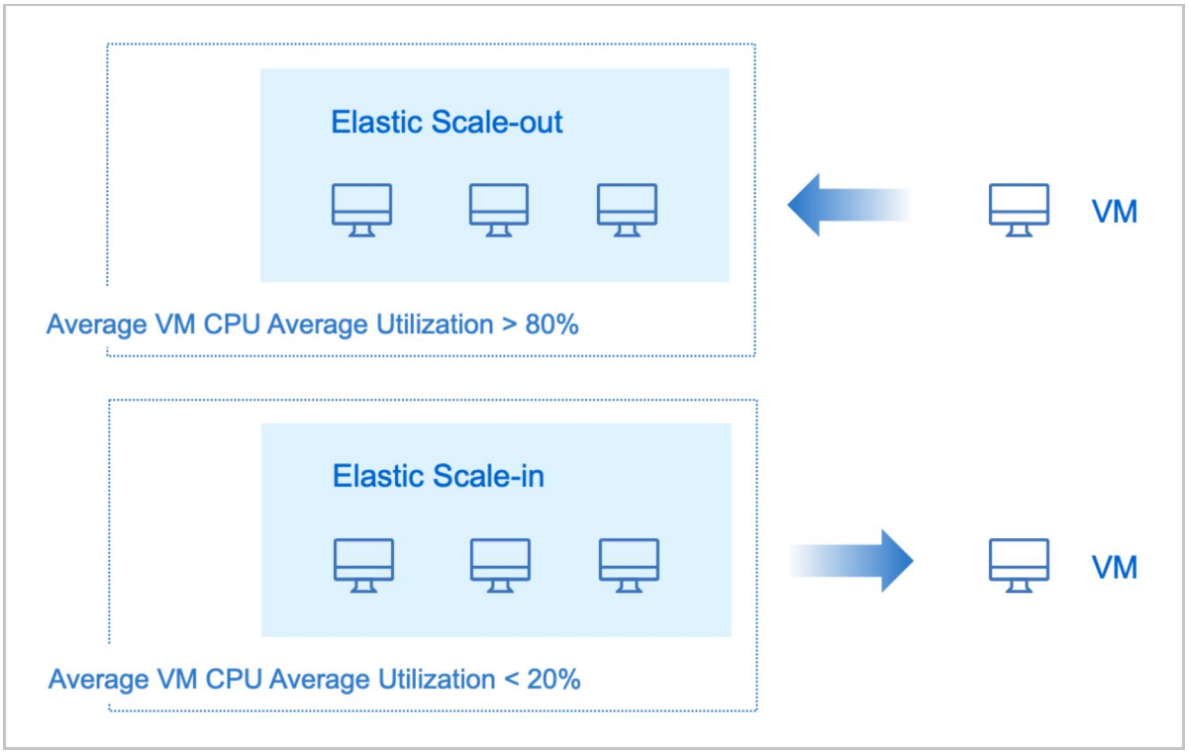

2. Elastic Self-Health

- In the elastic self-health mode, an auto-scaling group monitors the health status of the VM instances within the group, automatically removes the unhealthy VM instances, and adds additional VM instances. In this regard, healthy VM instances within the group will be ensured to be automatically adjusted not lower than the minimum specified number of VM instances.
- Two types of health check are provided to trigger elastic self-health, including load balancer health check and VM health check. If you configured the load balancing service for an auto-scaling group, we recommend that you select the health check mechanism native to a load balancer.

Figure 2-4: Elastic Self-Health

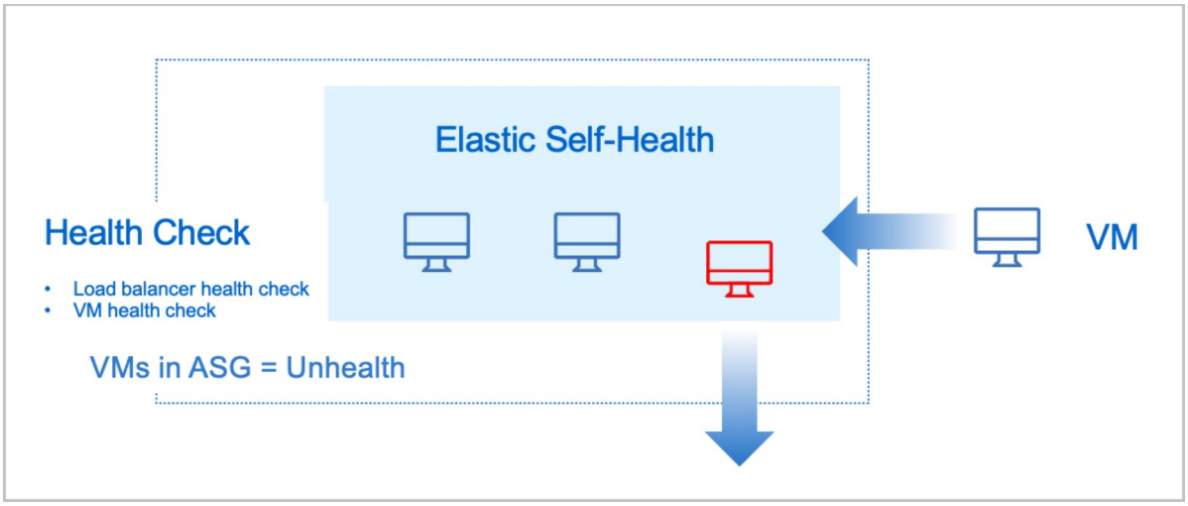

2.2.1.1.8 Snapshot

A snapshot is a point-in-time capture of data status in a volume. Before you perform a business-sensitive operation, you can schedule snapshot creation at specified time points to record the state of the root volume, data volume, or memory of a VM instance. This allows rollback in case of breakdowns. If you want to backup data for a long term, you can use the Backup Service.

2.2.1.1.9 VM Scheduling Policy

A VM scheduling policy is a resource orchestration policy based on which VM instances are assigned hosts to achieve the high performance and high availability of businesses.

Concepts

The VM scheduling policy involves the following key concepts:

- NexaVM Cloud provides four types of scheduling policies, and each policy can be executed based on two execution mechanisms. You can schedule a VM instance by associating it with different scheduling policies and execution mechanisms to meet the needs of different business scenarios.
  - Four types of VM scheduling policies: VM Exclusive from Each Other, VM Affinitive to Each Other, VMs Exclusive from Hosts, and VMs Affinitive to Hosts.
    - VM Exclusive from Each Other: VM instances in the same VM scheduling group should not/must not run on the same host.

- VM Affinitive to Each Other: VM instances in the same VM scheduling group should/must run on the same host.
- VMs Exclusive from Hosts: Given any one of the VM instances in a VM scheduling group and any one of the hosts in a host scheduling group, the VM instance should not/must not run the host.
- VMs Affinitive to Hosts: Given any one of the VM instances in a VM scheduling group and any one of the hosts in a host scheduling group, the VM instance should/must run the host.

— A VM scheduling policy can be executed based on either of the two execution mechanisms: Hard and Soft. Hard execution mechanism requires VM instances to strictly comply with their associated scheduling policies, while Soft execution mechanisms allows for some flexibility in policy execution when the hosts do not have sufficient resources.

- Hard: VM instances are forcibly assigned to 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 VM instances 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 VM instance, the VM instance will end up failure upon startup.
- Soft: VM instances are primarily assigned to 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 VM instances 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 VM instance, the VM instance will attempt to run on a host that does not satisfy the policy.

• A VM scheduling group is the basic unit for VM instances to associate with the VM scheduling policies.

— A VM instance can only be added to one VM scheduling group. After the addition, the VM instance will be scheduled based on the associated scheduling policies.

— A VM scheduling group can associate with one or more scheduling policies.

— A scheduling policy can associate with one VM scheduling group.

— Deleting a VM scheduling group also deletes its associated VM scheduling policies.

- Host scheduling group is the basic unit for hosts to associate with host scheduling policies. You can use a host scheduling group when you select **VMs Affinitive to Hosts** or **VMs Exclusive from Hosts** policies.

  - A host can only be added to one host scheduling group. After the addition, the host will be scheduled based on the associated scheduling policies.
  - A host scheduling group can associate with one or more scheduling policies.
  - A scheduling policy can associate with one VM scheduling group.
  - Deleting a host scheduling group also deletes its associated VM scheduling policies.

Fundamentals

NexaVM Cloud supports adding VM instances to a VM scheduling group, and executing a VM scheduling policy by associating a scheduling policy with the VM scheduling group.

- If you associate a VM scheduling group with a **VM Exclusive from Each Other** or **VM Affinitive to Each Other** scheduling policy, you do not need to specify a host scheduling group, as the VM instance will be assigned to hosts based on the policy and execution mechanism.
- If you associate a VM scheduling group with a **VMs Affinitive to Hosts** or **VMs Exclusive from Hosts** scheduling policy, you need to specify the corresponding host scheduling groups and the VM instance will be assigned to hosts based on the policy and execution mechanism.

In the following section, you will learn how the four scheduling policies work through four scenario illustrations.

Scenario 1: Assume that there are three hosts in the zone: Host1, Host2, and Host3. VM scheduling group A has been associated with the **VM Exclusive from Each Other** scheduling policy. VM1 and VM2 have joined VM scheduling group A and run on Host1 and Host2 respectively. In this setting, if you join VM3 to VM scheduling group A, the way how VM3 is scheduled under different execution mechanisms will be as follows:

- Under the Hard execution mechanism, VM3 follows the policy of **VM Exclusive from Each Other**:

  - If Host3 has sufficient resources, VM3 can start and run normally on Host3.
  - If Host3 does not have sufficient resources, then VM3 cannot start.
- Under the Soft execution mechanism, VM3 follows the policy of **VM Exclusive from Each Other** and first chooses to start on Host3:

- If Host3 has sufficient resources, VM3 can start and run normally on Host3.
- If Host3 does not have sufficient resources, VM3 tries to start on other host that has available resources. In this scenario, VM3 starts and runs on Host2.

Figure 2-5: VM Exclusive from Each Other (Hard/Soft)

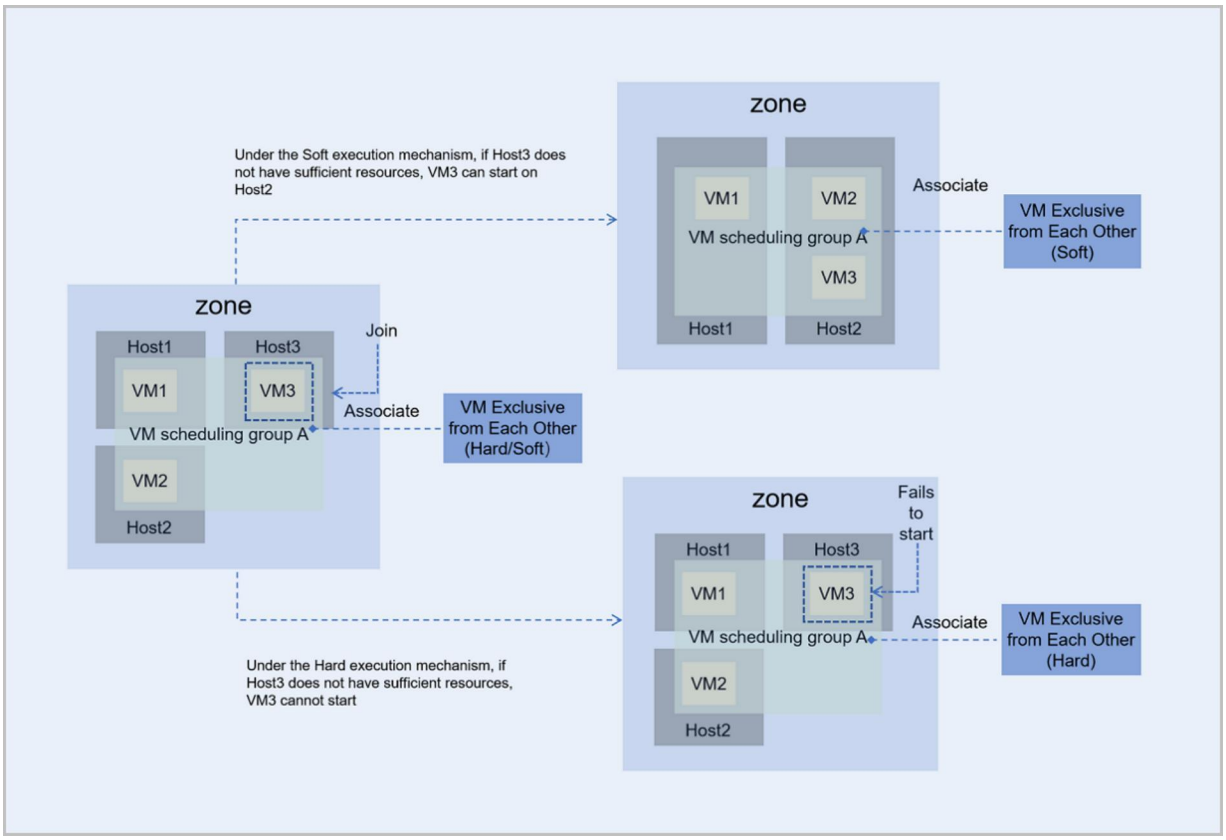

Scenario 2: Assume that there are two hosts in the zone: Host1 and Host2. VM scheduling group A has been associated with the **VM Affinitive to Each Other** scheduling policy. VM1 and VM2 have joined VM scheduling group A and run on Host1. In this setting, if you join VM3 to VM scheduling group A, the way how VM3 is scheduled under different execution mechanisms will be as follows:

- Under the Hard execution mechanism, VM3 follows the policy of **VM Affinitive to Each Other**:
  - If Host1 has sufficient resources, VM3 can start and run normally on Host1.
  - If Host1 does not have sufficient resources, VM3 cannot start.
- Under the Soft execution mechanism, VM3 follows the policy of **VM Affinitive to Each Other** and first chooses to start on Host1:
  - If Host1 has sufficient resources, VM3 can start and run normally on Host1.

- If Host3 does not have sufficient resources, VM3 tries to start on other host that has available resources. In this scenario, VM3 starts and runs on Host2.

Figure 2-6: VM Affinitive to Each Other (Hard/Soft)

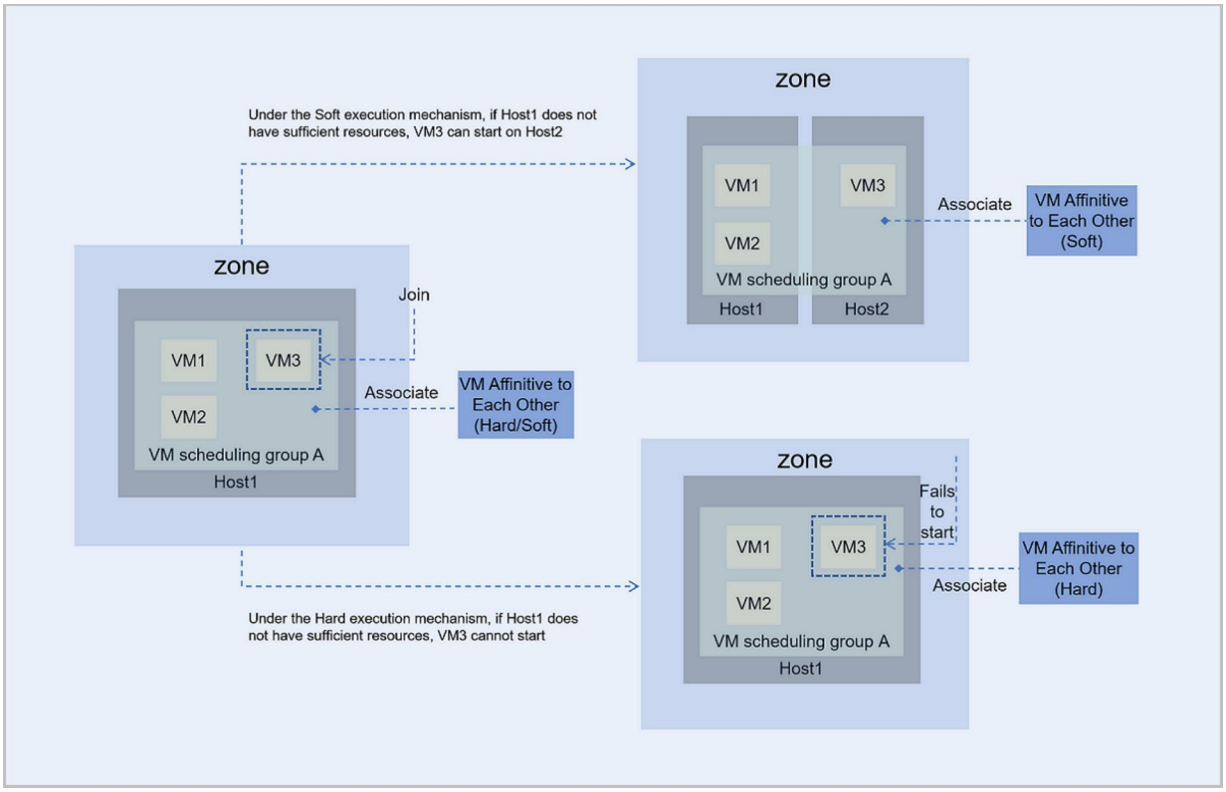

Scenario 3: Assume that there are three hosts in the zone: Host1, Host2, and Host3. VM scheduling group A has been associated with the **VMs Exclusive from Hosts** scheduling policy. VM1 and VM2 have joined VM scheduling group A and run on Host1. Host scheduling group A has also been associated with the **VMs Exclusive from Hosts** scheduling policy. Host2 and Host3 have joined host scheduling group A, with each of these hosts running two VMs respectively. In this setting, if you join VM3 to VM scheduling group A, the way how VM3 is scheduled under different execution mechanisms will be as follows:

- Under the Hard execution mechanism, VM3 and hosts in the host scheduling group A follow the policy of **VMs Exclusive from Hosts**:
  - If Host1 has sufficient resources, VM3 can start and run normally on Host1.
  - If Host1 does not have sufficient resources, VM3 cannot start.
- Under the Soft execution mechanism, VM3 and hosts in the host scheduling group A follow the policy of **VMs Exclusive from Hosts** and VM3 first chooses to start on Host1:
  - If Host1 has sufficient resources, VM3 can start and run normally on Host1.

— If Host1 does not have sufficient resources, VM3 tries to start on other host that has available resources. In this scenario, VM3 starts and runs on Host2.

Figure 2-7: VMs Exclusive from Hosts (Hard/Soft)

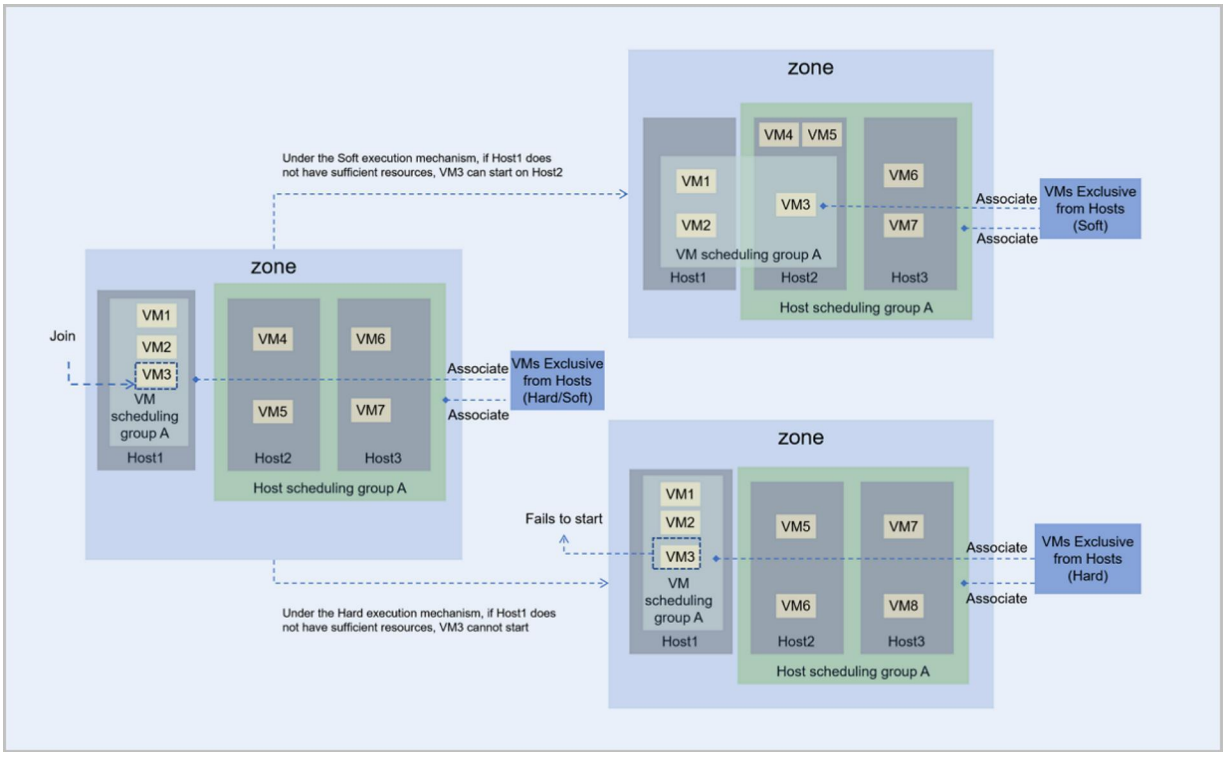

Scenario 4: Assume that there are three hosts in the zone: Host1, Host2, and Host3. VM scheduling group A has been associated with the **VMs Affinitive to Hosts** scheduling policy. VM1~VM5 have joined VM scheduling group A and run on Host2 and Host3 respectively. Host scheduling group A has also been associated with the **VMs Affinitive to Hosts** scheduling policy, and Host2 and Host3 have joined the host scheduling group A. In this setting, if you join VM6 to VM scheduling group A, the way how VM6 is scheduled under different execution mechanisms will be as follows:

- Under the Hard execution mechanism, VM6 and hosts in the host scheduling group A follow the policy of **VMs Affinitive to Hosts**:
  - If Host2 or Host3 have sufficient resources, VM6 can start and run normally on Host2 or Host3.
  - If Host2 and Host3 do not have sufficient resources, VM6 cannot start.
- Under the Soft execution mechanism, VM6 and hosts in the host scheduling group A follow the policy of **VMs Affinitive to Hosts** and VM6 first chooses to start on Host2 or Host3:

- If Host2 or Host3 have sufficient resources, VM6 can start and run normally on Host2 or Host3.
- If both Host2 and Host3 do not have sufficient resources, VM6 tries to start on other host that has available resources. In this scenario, VM6 starts and runs on Host1.

Figure 2-8: VMs Affinitive to Hosts (Hard/Soft)

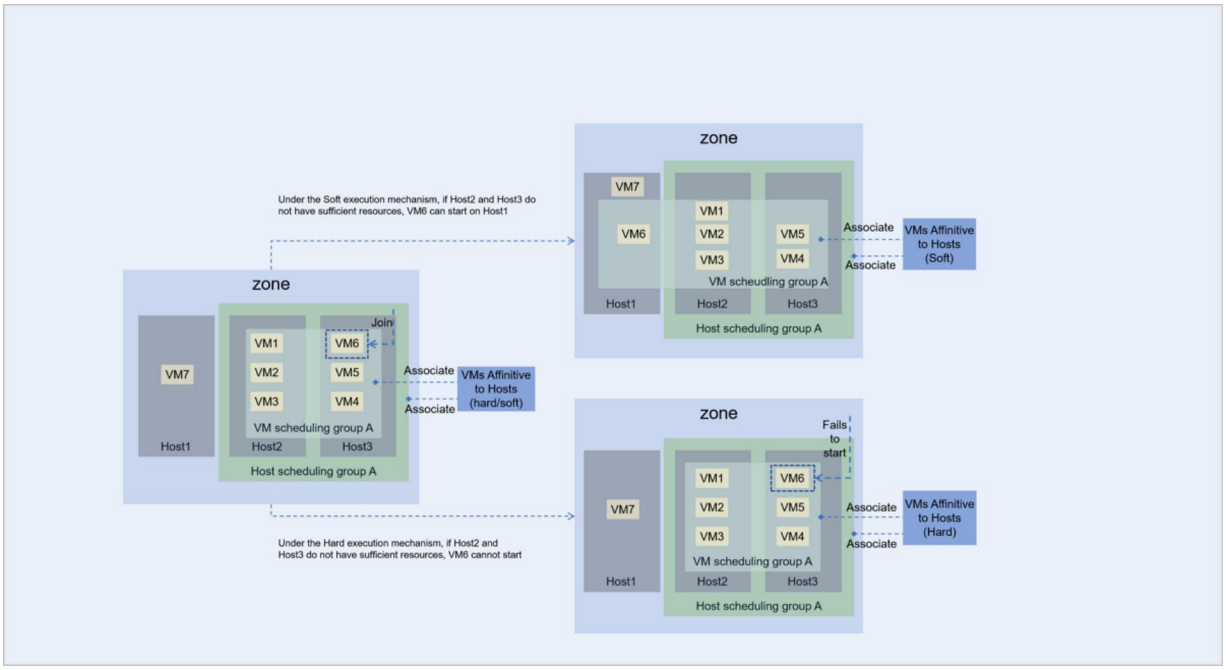

Advantages

VM scheduling policy has the following advantages:

- **Comprehensive & Flexible**:
  - NexaVM Cloud provides four types of scheduling policies matched with two execution mechanisms to define mutual exclusion/affinity relationships between VM instances and between VM instances and hosts. Various scheduling policies can be flexibly combined to meet the needs of all mainstream business scenarios.
  - NexaVM Cloud supports the association of one or more scheduling policies with multiple VM instances in the VM scheduling group, as well as the removal of these policies from the VM instances in the VM scheduling group, which is simple and efficient.
  - NexaVM Cloud intuitively displays the VM scheduling status and provides quick conflict resolution operations to facilitate users to grasp business scheduling dynamics in real time and make timely adjustments.
- **Powerful & Reliable**:

- NexaVM Cloud supports exclusion/affinity between the same/different businesses, achieving business isolation/efficient communication and ensuring high business performance and reliability.
- You can flexibly configure the failure domains in VM businesses through host scheduling groups. NexaVM Cloud supports single host deployment, batch hosts deployment within a single cluster, and cross-cluster hosts deployment to avoid single point of failure. This ensures business stability and improves physical resource utilization.

Use Cases

In the following part, we introduce some use cases for the **VM Exclusive from Each Other** (Hard/Soft) and **VMs Affinitive to Hosts** (Hard/Soft) scheduling policies.

- **VM Exclusive from Each Other** (Hard):

  In this case, we have two VM instances that run an active-backup database. The two VM instances are required to be deployed on different hosts to ensure business high availability.

  - Example: A user deploys two VMs for hosting a main and a backup MySQL database respectively, with a requirement that the two VM instances to be deployed on different hosts to minimize the risk of business downtime. Due to automatic deployment, the user does not know which hosts have available resources. At this point, the user can choose a **VM Exclusive from Each Other** (Hard) policy to ensure that the two VMs run on two different hosts, ensuring the business high availability.

- **VM Exclusive from Each Other** (Soft):

  In this case, we want the nodes with different roles in Hadoop to be spread across different hosts, so as to improve the overall system performance.

  - Example: When deploying a Hadoop system, the user cannot predict the total number of nodes for different roles such as namenode, datanode, jobtracker, tasktracker. However, it's clear that deploying these nodes on different hosts would enhance efficiency. The **VM Exclusive from Each Other** (Soft) policy can help distribute the Hadoop cluster across as many different hosts as possible, thereby alleviating I/O pressure and improving the overall system performance.

- **VMs Affinitive to Hosts** (Hard):

  In this case, we want to deploy the business VM instances on hosts with a specified CPU frequency, thereby ensuring the business stability.

- Example: A user deploys four VM instances running important businesses and requires these VM intances run on the hosts that have a high CPU frequency. Given that there are limited hosts that can meet this CPU frequency requirement, the user can choose a **VMs Affinitive to Hosts** (Hard) policy to force these VM instances run on the specified hosts to ensure business stability.

- **VMs Affinitive to Hosts** (Soft):

  In this case, we want VM instances that run different businesses are deployed to hosts in the same rack as much as possible, so as to facilitate efficient communication between businesses.

  - Example: A user deploys four VM instances that run different businesses that require frequent intercommunication. In this case, minimizing the physical distance between the VM instances can significantly reduce communication latency. At this point, the user can choose a **VMs Affinitive to Hosts** (Soft) policy to deploy the VM instances on the specified hosts, thereby promoting efficient business intercommunication.

2.2.1.2 Hardware

2.2.1.2.1 Zone

A zone is a logical group of resources such as clusters, L2 networks, and primary storage. Zone is the largest resource scope defined in the Cloud.

- In a data center, a zone corresponds to an equipment room.
- A zone defines a visible boundary. Sub-resources within the same zone are visible mutually and can form a certain relationship. However, sub-resources within different zones are invisible mutually and cannot form mutual relationships.

2.2.1.2.2 Cluster

A cluster is a logical group of hosts (compute nodes). In a real data center, a cluster usually maps to a rack.

When you plan a cluster, note that:

- All hosts in the same cluster must be installed with the same operating system.
- All hosts in the same cluster must have the same network configuration.
- All hosts in the same cluster must be able to access the same primary storage.

- Before a cluster can provide VM services, the cluster must have a primary storage and an L2 network attached.
- The scale of a cluster, which is the maximum number of hosts that the cluster can contain, is not limited.

The relationship between a typical cluster and its associated resources is as follows:

Cluster | Zone

You can create more than one cluster in a zone, and allocate newly created hosts to different clusters as needed.

Cluster | Primary Storage and L2 Network

You can attach primary storage and L2 networks to or detach them from a cluster. The following diagram shows the relationship between cluster and primary storage, L2 network.

Figure 2-9: Relationship Between Cluster and Primary Storage, L2 Network

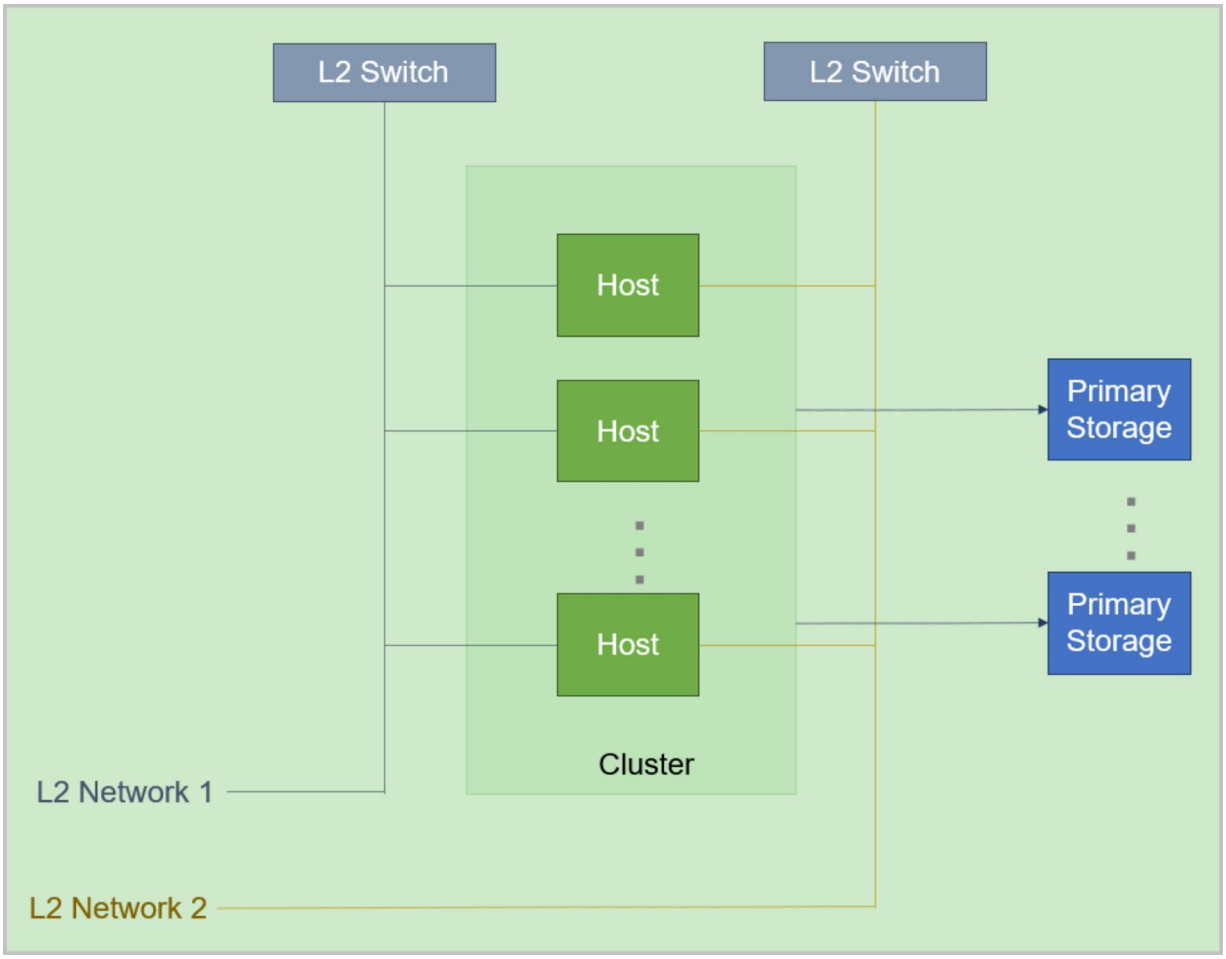

Note:

When you attach a primary storage and an L2 network to a cluster, note that:

1. Cluster | Primary Storage

- A primary storage can be attached to one or more clusters.
- A cluster can have one or more primary storage attached.

  The following are primary storage of the same type that a cluster can have:

  - A cluster can have one or more LocalStorage primary storage attached.
  - A cluster can have one or more NFS primary storage attached.
  - A cluster can have one or more SharedBlock primary storage attached.
  - A cluster can have one SharedMountPoint (SMP) primary storage attached.
  - A cluster can have only one Ceph primary storage attached.
  - A cluster can have only one AliyunNAS primary storage attached.
  - A cluster can have only one AliyunEBS primary storage attached.

  The following are combinations of primary storages that a cluster can have:

  - A cluster can have both a LocalStorage and an NFS primary storage attached.
  - A cluster can have both a LocalStorage and an SMP primary storage attached.
  - A cluster can have both a LocalStorage and a SharedBlock primary storage attached.
  - A cluster can have both a Ceph and a maximum of three LocalStorage primary storage attached.
  - A cluster can have both a Ceph and a SharedBlock primary storage attached.
  - A cluster can have both a Ceph and more than one SharedBlock primary storage attached.

  The following table lists the relationship between a primary storage and a cluster.

  Table 2-1: Relationship Between Primary Storage and Cluster

Primary Storage	Cluster
LocalStorage	A cluster can have one or more LocalStorage primary storage attached.
NFS	A cluster can have one or more NFS primary storage attached.

Primary Storage	Cluster
SharedBlock	A cluster can have one or more SharedBlock primary storage attached.
SMP	A cluster can have one SMP primary storage attached.
Ceph	A cluster can have only one Ceph primary storage attached.
AliyunNAS	A cluster can have only one AliyunNAS primary storage attached.
AliyunEBS	A cluster can have only one AliyunEBS primary storage attached.
LocalStorage + NFS	A cluster can have one LocalStorage and one NFS primary storage attached.
LocalStorage + SMP	A cluster can have one LocalStorage and one SMP primary storage attached.
LocalStorage + SharedBlock	A cluster can have one LocalStorage and one SharedBlock primary storage attached.
Ceph + LocalStorage	A cluster can have both a Ceph and a maximum of three LocalStorage primary storage attached.
Ceph + SharedBlock	• A cluster can have one Ceph and one SharedBlock primary storage attached. • A cluster can have one Ceph and multiple SharedBlock primary storage attached.

- When you attach multiple LocalStorage primary storage to a cluster, partition the corresponding URLs on the hosts before you add hosts and primary storage, and make sure that each LocalStorage is deployed on an exclusive logical volume or physical disk.
- A primary storage can be accessed by all hosts in the cluster to which the primary storage belongs.
- If a primary storage cannot be accessed by hosts in the cluster due to network typology changes in the data center, you can detach the primary storage from the cluster.

2. Cluster | L2 Network

- A cluster can have one or more L2 networks attached. Also, an L2 network can be attached to one or more clusters.

- A cluster can have a VXLAN pool attached. The VNIs in the VXLAN pool can be used to create different VXLAN networks.
- One NIC can be used to create only one NoVlan network.
- For VLAN networks, different VLAN IDs represent different L2 networks.
- If hosts in a cluster no longer exist in the layer 2 broadcast domain of an L2 network due to network typology changes in the data center, you can detach the L2 network from the cluster.

Cluster | Backup Storage

No direct dependency exists between a cluster and a backup storage. A backup storage can provide services for multiple clusters.

The following table lists the relationship between primary storage (PS) and backup storage (BS).

Table 2-2: Relationship Between PS and BS

PS/BS	ImageStore	Ceph
LocalStorage	○	×
NFS	○	×
SMP	○	×
Ceph	○	○
SharedBlock	○	×

- When primary storage are LocalStorage, NFS, or SMP, the default type for backup storage is ImageStore.
- When primary storage are NFS or SMP, you can manually mount the corresponding shared directories to the local directories of the corresponding backup storage. In this regard, both primary storage and backup storage can use the network shared storage.
- When primary storage is Ceph, you can use the primary storage in the same Ceph cluster as backup storage. You can also use the ImageStore primary storage as backup storage.
- When primary storage is SharedBlock, the default type for backup storage is ImageStore.
- When primary storage is AliyunNAS, the default type for backup storage is ImageStore.
- When primary storage is AliyunEBS, the default type for backup storage is AliyunEBS.

2.2.1.2.3 Host

A host provides compute, network, and storage resources for VM instances. Hosts are core resources on the Cloud. VM instances are running on hosts.

Figure 2-10: Host

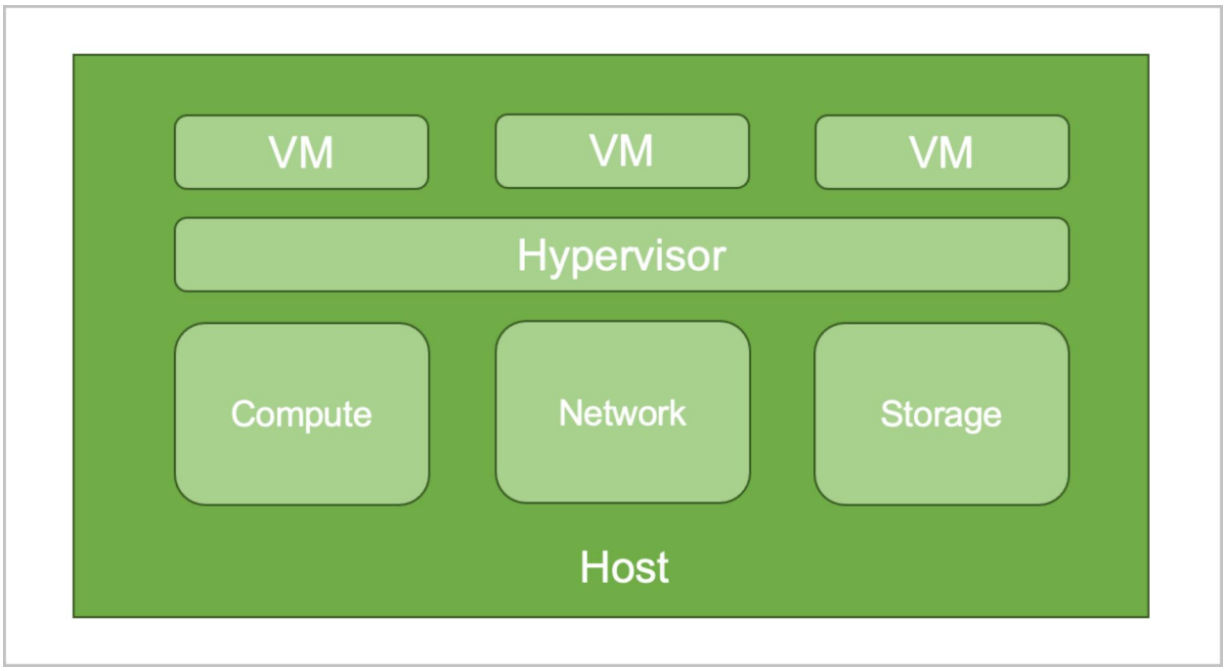

2.2.1.2.4 Primary Storage

A primary storage is one or more servers that store volume files of VM instances. These files include root volume snapshots, data volume snapshots, image caches, root volumes, and data volumes.

The Cloud supports the following types of primary storage:

- **Local Storage:** This type of primary storage uses the hard disks to store disk files.
- **Network Shared Storage:** This type of primary storage supports NFS, Shared Mount Point (SMP), Ceph, SharedBlock, AliyunNAS, and AliyunEBS.
  - NFS primary storage uses the Network File System (NFS) to store files.
  - SMP primary storage supports network shared storage provided by commonly used distributed file systems, such as MooseFS, GlusterFS, OCFS2, and GFS2.
  - Ceph primary storage uses distributed block storage to store files.
  - SharedBlock primary storage uses shared block storage to store files.
  - AliyunNAS primary storage uses distributed files to store files.

— AliyunEBS primary storage uses distributed block storage to store files.

The following table lists the relationship between primary storage and a cluster.

Table 2-3: Relationship Between Primary Storage and Cluster

Primary Storage	Cluster
LocalStorage	A cluster can have one or more LocalStorage primary storage attached.
NFS	A cluster can have one or more NFS primary storage attached.
SharedBlock	A cluster can have one or more SharedBlock primary storage attached.
SMP	A cluster can have one SMP primary storage attached.
Ceph	A cluster can have only one Ceph primary storage attached.
AliyunNAS	A cluster can have only one AliyunNAS primary storage attached.
AliyunEBS	A cluster can have only one AliyunEBS primary storage attached.
LocalStorage + NFS	A cluster can have one LocalStorage and one NFS primary storage attached.
LocalStorage + SMP	A cluster can have one LocalStorage and one SMP primary storage attached.
LocalStorage + SharedBlock	A cluster can have one LocalStorage and one SharedBlock primary storage attached.
Ceph + LocalStorage	A cluster can have both a Ceph and a maximum of three LocalStorage primary storage attached.
Ceph + SharedBlock	• A cluster can have one Ceph and one SharedBlock primary storage attached. • A cluster can have one Ceph and multiple SharedBlock primary storage attached.

2.2.1.2.5 Backup Storage

A backup storage is a storage server that stores VM image templates, including ISO image files.

- A backup storage must be attached to a zone before the resources in the zone can access it. Note that you can share images across multiple zones by using the backup storage.

Figure 2-11: Backup Storage

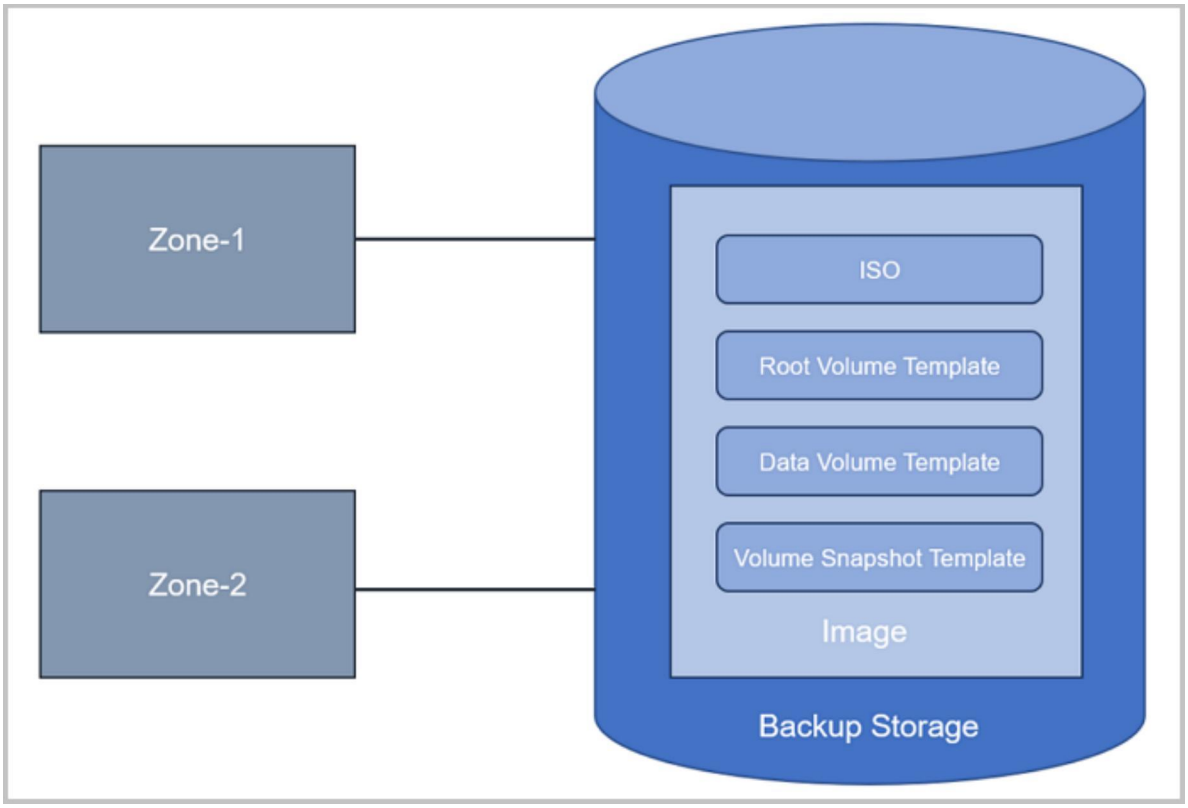

- To better manage backup storage and zones, the UI specifies that one backup storage can only correspond to one zone. In the UI, when you add a backup storage, the backup storage will be attached to the current zone by default. When you delete a zone, the backup storage attached to the zone will also be deleted.

Backup Storage Type

The Cloud supports the following types of backup storage:

1. **ImageStore**
   - Stores image files by means of image slices and supports incremental storage.
   - Allows you to create snapshots and images when VM instances are running or stopped.
   - Allows you to clone VM instances without data volumes when these VM instances are running, paused, or stopped.

- Allows you to clone VM instances with data volumes when these VM instances are running, paused, or stopped, and with storage types of LocalStorage, NFS, SharedMountPoint (SMP), Ceph, or SharedBlock.
- Supports image synchronization across ImageStore backup storage on the same management node.
- Allows you to obtain the existing image files under the URL path in the backup storage.

2. **Ceph**

- Stores image files by means of Ceph distributed block storage.
- Allows you to create snapshots and images when VM instances are running or stopped.
- Allows you to clone VM instances without data volumes when these VM instances are running, paused, or stopped.
- Allows you to clone VM instances with data volumes when these VM instances are running, paused, or stopped, and with the storage type of Ceph.
- Allows you to export images on the UI or backup storage.
  - You can export images, copy exported image URLs, and download exported images on the UI.
  - You can also export images on a backup storage.

    For example, assume that the image path you use is *ceph://bak-t-c9923f9821bf45498fdf 9cdfa1749943/61ece0adc7244b0cbd12dafbc5494f0c*.

    Then, run the following command on the backup storage:

```
rbd export -p bak-t-c9923f9821bf45498fdf9cdfa1749943 --image
61ece0adc7244b0cbd12dafbc5494f0c --path /root/export-test.image

# bak-t-c9923f9821bf45498fdf9cdfa1749943 is the name of the
pool where the image resides.
# 61ece0adc7244b0cbd12dafbc5494f0c is the name of the image.
# /root/export-test.image is the name of the exported file.
```

3. **AliyunEBS**

- Stores image by means of object storage.
- Allows you to create snapshots and images when VM instances are running or stopped.
- Allows you to clone VM instances without data volumes when these VM instances are running, paused, or stopped.
- Does not allow you to clone VM instances with data volumes.

- Allows you to export images on backup storage. For more information, contact the official technical support.

Backup Storage | Primary Storage

The following table lists the relationship between primary storage (PS) and backup storage (BS).

Table 2-4: Relationship Between PS and BS

PS/BS	ImageStore	Ceph
LocalStorage	○	×
NFS	○	×
SMP	○	×
Ceph	○	○
SharedBlock	○	×

- When primary storage are LocalStorage, NFS, or SMP, the default type for backup storage is ImageStore.
- When primary storage are NFS or SMP, you can manually mount the corresponding shared directories to the local directories of the corresponding backup storage. In this regard, both primary storage and backup storage can use the network shared storage.
- When primary storage is Ceph, you can use the primary storage in the same Ceph cluster as backup storage. You can also use the ImageStore primary storage as backup storage.
- When primary storage is SharedBlock, the default type for backup storage is ImageStore.
- When primary storage is AliyunNAS, the default type for backup storage is ImageStore.
- When primary storage is AliyunEBS, the default type for backup storage is AliyunEBS.

2.2.1.2.6 SAN Storage

Storage Area Network (SAN) storage can be categorized into iSCSI storage and FC storage:

- iSCSI storage is an SAN storage that uses the iSCSI protocol for data transmission. You can add an iSCSI SAN block as a Shared Block primary storage or pass through the block to a VM instance.
- FC storage is an SAN storage that uses the FC technology for data transmission. You can add an FC SAN block as a Shared Block primary storage or pass through the block to a VM instance. NexaVM Cloud supports FC for SAN storage connection with multi-path I/O.

Scenarios

- Passes through an iSCSI or FC disk to a VM instance.
- Adds an iSCSI or FC disk as a shared block and takes the shared block as a Shared Block primary storage.

Considerations

iSCSI storage:

- You can add a disk that is not attached to a VM instance as a Shared Block primary storage.
- A logical unit number (LUN) that is added as a primary storage cannot be used for other purposes.
- You can attach a disk that is not added as a primary storage to a VM instance.
- You can attach a disk to multiple VM instances or attach multiple disks to one VM instance.

FC storage:

- If a block device is not attached to a VM instance and the cluster where the block device resides is in normal state, you can add the block device as a Shared Block primary storage.
- A logical unit number (LUN) that is added as a primary storage cannot be used for other purposes.
- You can attach a block device that is not added as a primary storage to a VM instance.
- You can attach a block device to multiple VM instances or attach multiple block devices to one VM instance.

2.2.1.2.7 Physical Network

Physical Network: Real networks deployed in a physical environment, which can be divided into different categories based on their application scenarios, such as management networks, storage networks, business networks, backup networks, and migration networks. You can attach tags to a NIC port on NexaVM Cloud to mark which type of network it belongs to according to your physical environment planning.

Physical Network Type

Physical networks can be divided into the following categories:

- **Management Network**: A network used to manage physical resources on NexaVM Cloud. The management nodes communicate with the hosts, primary storage, backup storage, VPC vRouters, and LB instances (performance-dedicated) on the Cloud through the management

network. If you deploy a management network, it can be automatically identified by NexaVM Cloud.

- The management networks you add in **Resource Pool** > **Network Resource** > **Dedicated Network** comply with the management network definition here.
- The host IP addresses displayed on NexaVM Cloud belong to the management network.
- The host callback IP addresses belong to the management network.

- **Storage Network**: A network specified for storage services. NexaVM Cloud uses a storage network to get the health status of VM instances. We recommend that you deploy a dedicated storage network to avoid possible risks.
- **Business Network**: A network that can be added as an L2/L3 network on and attached to VM instances on NexaVM Cloud. A VM instance uses a business network to provide services.
- **Backup Network**: A network used to back up or recover local business data in a local backup scenario. You need to deploy a backup network if you use Backup Service.
- **Migration Network**: A network used for VM migrations on NexaVM Cloud.

Figure 2-12: Physical Network Deployment

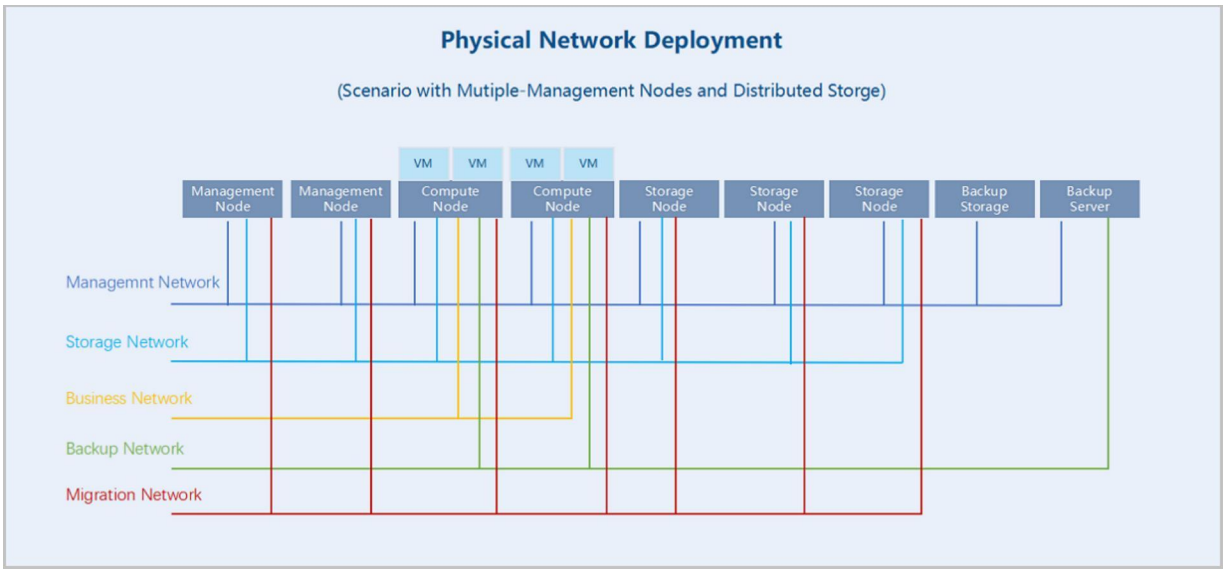

Characteristics

- The **Physical Network** page displays physical NIC ports (bonds and non-bonds) in fact. These NIC ports realize the network connections between hosts. You manage physical networks by managing these NIC ports.

- You can attach network-type tags to NIC ports according to your physical environment planning. Management network tags are attached to corresponding ports automatically by the system and other tags need to be attached by yourself.
  - If you add a management network in **Resource Pool** > **Network Resource** > **Dedicated Network**, the NIC port of the L2 network it belongs to is attached with a management network tag.
  - The NIC port of the network to which a host IP address belongs is attached with a management network tag.
  - The NIC port of the network to which a host callback IP address belongs is attached with a management network tag.
- After you attach tags to corresponding NIC ports, you can view the flow monitoring of different network types, which help you master accurate network workloads more flexibly.

2.2.1.3 Network Resource

2.2.1.3.1 SDN Controller

An SDN controller is used to control network devices such as switches. You can add an external SDN controller to the Cloud and use the controller to control external switches and other network devices. By adding an SDN controller, you can take over SDN networks of hardware switches to reduce network latencies and improve VXLAN network performances.

- Before you can add an SDN controller to the Cloud, you need to plan management networks in advance, and complete the basic configurations of the SDN controller.
- Currently, the Cloud supports only the H3C VCFC SDN controller.

Note:

If you use VCFC to configure hardware SDN, configure the mapping between VLAN and VXLAN on VCFC in advance.

2.2.1.3.2 L2 Network Resources

VXLAN Pool

VXLAN Pool: A VXLAN pool is a collection of VXLAN networks established based on VXLAN Tunnel Endpoints (VTEPs). The VNI of each VXLAN network in a VXLAN pool must be unique.

- Before you can use a VXLAN network, create a VXLAN pool in advance.

- A VXLAN pool is only a collection of VXLAN networks and cannot be used to create L3 networks.
- VXLAN pool supports two types of SDN: software SDN and hardware SDN.
  - Software SDN:
    - The VNI range of a software-SDN VXLAN pool can be 1-16777214.
    - Hosts in the cluster to which the software-SDN VXLAN pool is attached must have an IP address as a VXLAN tunnel endpoint (VTEP) in the specified CIDR.
      - Generally, a VTEP corresponds to a NIC IP of a compute node within a cluster. On the Cloud, you can configure a VTEP according to its CIDR. For example,
        - Assume that the NIC IP of a compute node is *10.12.0.8*, the netmask is *255.0.0.0*, and the gateway is *10.0.0.1*. Then, the CIDR of the VTEP is *10.0.0.1/8*.
        - Assume that the NIC IP of the compute node is *172.20.12.13*, the netmask is *255.255.0.0*, and the gateway is *172.20.0.1*. Then, the CIDR of the VTEP is *172.20.0.1/16*.
      - When a VXLAN pool is attached to a cluster, the IP address that is associated to the VTEP will be checked without checking physical L2 devices.
  - Hardware SDN:
    - Before you can create a hardware-SDN VXLAN pool, add an SDN controller to the Cloud in advance.
    - The VNI range of a hardware-SDN VXLAN pool depends on the distributed vSwitch to which an SDN controller corresponds.
    - The host NIC in the cluster to which the hardware-SDN VXLAN pool is attached must connect to a switch managed by the SDN controller.
    - If you use VCFC to configure hardware SDN, configure the mapping between VLAN and VXLAN on VCFC in advance.

L2 Network

An L2 network is a layer 2 broadcast domain used for layer 2 isolation. Generally, L2 networks are identified by names of devices on the physical network.

- VLAN, VXLAN, and SDN can be used as an L2 network.
- An L2 network is used to provide layer 2 isolation for an L3 network.

The following four types of L2 networks are supported:

1. **L2NoVlanNetwork**

   - You must specify the NIC name of the host.
   - The hosts in the cluster to which the L2 network is attached must have NICs that share the same name.
   - When the data packets of VM instances flow out of the host NIC and reach the physical switch, the data packets are not flagged with VLAN tags. Note that the physical switch must be in VLAN Access mode.
   - If you create an L2 network of the L2NoVlanNetwork type, a network bridge is created based on the specified host NIC.

2. **L2VlanNetwork**

   - You must specify the host NIC name and VLAN ID. NexaVM Cloud allows you to create virtual NICs and virtual switches, and supports IEEE 802.1Q VLAN trunking.
   - The hosts in the cluster to which the L2 network is attached must have NICs that share the same name.
   - When the data packets of VM instances flow out of the host NIC and reach the physical switch, the data packets are tagged with the specified VLAN ID.
   - If you create an L2 network of the L2VlanNetwork type, a VLAN device is created based on the specified VLAN ID and then a network bridge is created based on the VLAN device.
   - If you attach an L2 network of the L2NoVlanNetwork type and an L2 network of the L2VlanNetwork type or attach multiple L2 networks of the L2VlanNetwork type to a cluster , the physical switch NIC must be in Trunk mode. Besides, the VLAN ID in use must be contained in the Trunk VLAN configurations.

3. **VxlanNetwork**

   - Virtual Extensible LAN (VXLAN) is an overlay technology that allows for the creation of overlaying L2 networks. This technology can support a maximum of 16 million logical networks.
   - VxlanNetwork is an implementation of the software-based VXLAN technology.
   - If you create an L2 network of the VxlanNetwork type, you must specify a software SDN-based VXLAN pool. The L2 network must correspond to a VNI in the pool.
   - The VTEP IPs of the hosts in the cluster to which the L2 network of the VxlanNetwork type is attached must belong to the specified VXLAN pool.
   - When the data packets of VM instances flow out of the host, the host encapsulates VXLAN messages to the data packets and then sends the encapsulated data to the physical switch.

4. **HardwareVxlanNetwork**

   - Virtual Extensible LAN (VXLAN) is an overlay technology that allows for the creation of overlaying L2 networks. This technology can support a maximum of 16 million logical networks.
   - HardwareVxlanNetwork is a solution to the integration with third-party hardware SDN.
   - If you create an L2 network of the HardwareVxlanNetwork type, you must specify a hardware SDN-based VXLAN pool. The L2 network must correspond to a VNI in the pool.
   - When the data packets of VM instances flow out of the host NICs and reach the distributed virtual switches, the data packets are flagged with the specified VLAN ID. The VLAN ID is mapped with VXLAN ID based on the SDN controlled that you add to the Cloud.

Note:

- For some OS, the NIC name in the ethX format will be changed after the system reboots. In addition, the NIC sequence will also be randomly changed. We recommend that you change the NIC name of each compute node (especially for VM instances with multiple NICs) to a non-ethX format, such as em01.

2.2.1.3.3 L3 Network

An L3 network is a collection of network configurations for VM instances, including the network range, gateway, DNS, and network services.

- A network range includes an IP range (start IP and end IP), netmask, and gateway. For example, you can specify the IP range from *172.20.12.2* to *172.20.12.255*, set the netmask to *255.255.0.0*, and set the gateway to *172.20.0.1*. In addition, you can use a CIDR to specify a network range, such as *192.168.1.0/24*.
- DNS provides DNS resolution services used for configuring VM networks.

Concepts

- Public network: Generally, a public network is a logical network that is connected to the Internet. However, in an environment that has no access to the Internet, you can also create a public network.
  - A public network can be used in the flat network environment to create VM instances.
  - A public network can be used in the VPC network environment to create VM instances that work with public networks.

- Flat network: A flat network is connected to the network where the host is located and has direct access to the Internet. VM instances in a flat network can access public networks by using elastic IP addresses.
  - A flat network supports multiple network services, including DHCP, User Data, EIP, security group, and port mirroring.
  - The network services provided by a flat network use the distributed DHCP and the distributed EIP structure.
  - The DHCP service provided by a flat network also includes the DNS feature.
  - The network model used in the wizard is a flat network.
  - The flat network architecture based on VxlanNetwork or HardwareVxlanNetwork is supported.
- VPC network: A VPC network is a private network where VM instances can be created. A VM instance in a VPC network can access the Internet through a VPC vRouter.
  - A VPC network provides the following network services: DHCP, User Data, DNS, SNAT, route table, EIP, port forwarding, load balancing, IPsec tunnel, security group, dynamic routing, multicast routing, VPC firewall, port mirroring, and netflow.
  - The DHCP service of the VPC network uses DHCP by default.
  - VPC networks mainly use custom Linux VM instances as VPC vRouters to provide network services.
  - Network services can act on multiple subnets of a VPC at the same time, further improving network efficiencies.
  - Supports VxlanNetwork-based VPC network architecture.
  - Supports distributed routing, optimizing east-west network traffic and effectively reducing network latency.
- Dedicated network:
  - Management network: A management network is used to manage physical resources in the Cloud. For example, you can create a management network to manage access to hosts, primary storages, backup storages, and VPC vRouters.

Note:

When you create a VPC vRouter, you need an IP address that can be interconnected between the management nodes of the VPC vRouter. With this IP address, you can deploy an agent and obtain messages returned by the agent.

- Flow network: A flow network is a dedicated network for port mirror transmission. You can use a flow network to transmit the mirrors of data packets of NIC ports to the target ports. A flow network cannot be used for other purposes, such as creating VM instances.

- Specific network scenarios:
  - Storage network: A storage network is the network specified by the shared storage. You can use a storage network to check the health state of a VM instance. We recommend that you plan for an independent storage network in advance to avoid potential risks.
  - VDI network: When you create a cluster, you can specify CIDR for the VDI network in the cluster. In the VDI scenario, the network traffics generated by the protocol communicat ion between the server side and client side use the VDI network. If you do not make any configuration to the VDI network, notice that the management network will be used by default.
  - Migration network: When you create a cluster, you can specify CIDR for the migration network in the cluster. The migration network is used to migrate VM instances in the Cloud. If you do not make any configuration to the migration network, notice that the management network will be used for VM migrations.
  - Image synchronization network: An image synchronization network is used to synchronize images among ImageStore backup storages in the same management node.
    - If you deployed an independent network for synchronizing images, you can specify CIDR for the image synchronization network when you add an ImageStore backup storage.
    - If you do not make any configuration to the image synchronization network, notice that the management network will be used by default.
    - If you set an image synchronization network for both the source ImageStore backup storage and target one, only the image synchronization network in the target ImageStore backup storage takes effect.
  - Data network: A data network is the network where data can transfer between a compute node and a backup storage.
    - Using an independent data network can avoid network congestion and improve the data transfer rate.
    - If you do not make any configuration to the data network, notice that the management network will be used by default.

— Backup network: If you are using the Backup Service or the Continuous Data Protection (CDP) service, in the local backup scenario,both the data backup and recovery are implemented by using the backup network.

  - If you deploy an independent network for local backups, you can specify CIDR for the backup network when you add a local backup server.
  - Using an independent backup network can avoid network congestion and improve the data transfer rate.
  - If you do not make any configuration to the backup network, notice that the management network will be used for local backup by default.

Note:

The Backup Service and the CDP Service are separately provided in a separate module . To use this feature, purchase both the Base License and the Plus License. Note that a Base License is required before you can install a Plus License.

Considerations

- When you create a VM instance, you can specify multiple L3 networks, including flat networks, VPC networks, or a combination of flat networks and VPC networks.
- The Cloud supports multi-layer networks. In addition, the L2 networks of multi-layer networks can intercommunicate. Therefore, you need to pay a special attention to avoid the conflict of IP address spaces.
- You can use an L2 network to create multiple L3 networks. However, we recommend that unless necessary you do not create multiple L3 networks from an L2 network. This may cause the DHCP services of these L3 networks unable to work as expected.
- The network services and features supported by an L3 network are related to the network architecture model (flat network, VPC network) and the configured network protocol version (IPv4, IPv6). If a network is configured with both IPv4 and IPv6 protocols, servers of these two protocol types are loaded at the same time to provide corresponding services.

	IPv4	**IPv6**
Flat network	Supported network services: DHCP , User Data, EIP, security group, and port mirroring	Supported network services: DHCP , DNS, EIP, and security group
VPC network	Supported network services: DHCP , User Data, DNS, SNAT, route	Supported network services: DHCP , DNS, and security group

	IPv4	IPv6
	table, EIP, port forwarding, load balancing, IPsec tunnel, security group, dynamic routing, multicast routing, VPC firewall, port mirroring, and Netflow	
	Supported network service: VPC vRouter HA group	Supported network service: VPC vRouter HA group

- If you use an L2 network of the HardwareVxlanNetwork type, the L3 network created from the L2 network supports only flat network and corresponding network services.
- If you use the Smart NIC network acceleration mode for an L2 network, you can use the L2 network to create only IPv4 VPC networks and provide corresponding network services (excluding security group and port mirroring services).
- DHCP: By default, the VPC network provides distributed DHCP services by using the flat network service module.
- DNS: A VPC vRouter can act as a DNS server to provide DNS services. The DNS address in a VM instance is the IP address of the VPC vRouter. Note that the DNS address that you set is forwarded by the VPC vRouter.
- SNAT: A VPC vRouter can provide the source network address translation (SNAT) services for VM instances. Then, the VM instances can directly access the Internet by using SNAT.
- Route table: You can manage and customize routes through route tables.
- Security group: The security group service is provided by the security group network service module. You can configure and manage firewalls for VM instances by using iptables.
- Elastic IP address (EIP): You can bind an EIP to a VPC network. Then, the public network can interconnect with the private network of the VM instance.
- Port forwarding: The port forwarding service allows a public IP address to interconnect with the private IP address of a VM instance. To be more specific, you can create port forwarding rules to allow external networks to reach specific ports of your VM instances.
- Load balancing: The load balancing service distributes your inbound traffics from a public IP address to a group of backend VM instances. Then, this service automatically checks and isolates the VM instances that are unavailable.
- IPsec tunnel: The IPsec tunnel can be used to achieve interconnection between different virtual private networks (VPNs).

- Dynamic routing: The VPC vRouter supports the Open Shortest Path First (OSPF) routing protocol, which is used to distribute routing information within a single autonomous system.
- Multicast routing: The VPC vRouter forwards the multicast information sent by the multicast source to VM instances, achieving one-to-multi-point communication in the transmission side and receiving side.
- VPC firewall: The VPC firewall filters the south-north traffic on the VPC vRouter ports, effectivel y protecting the VPC communication security and VPC vRouter security.
- Netflow: The Netflow service monitors and analyzes the inbound and outbound traffics of the VPC vRouter NICs. Currently, the following two types of data-flow output format are supported: Netflow V5 and Netflow V9.

2.2.1.3.4 VPC

A Virtual Private Cloud (VPC) is a custom network environment that consists of VPC vRouters and VPC networks. With VPC, enterprise users can build a logically isolated private cloud.

VPC vRouter and VPC Network

A VPC consists of VPC vRouters and VPC networks.

- A VPC vRouter is a virtual router created from a vRouter offering. By default, a VPC vRouter has two types of network: public network and management network.
- A VPC network can be used as a VPC private network, and can be attached to a VPC vRouter.

Figure 2-13: VPC Network Topology

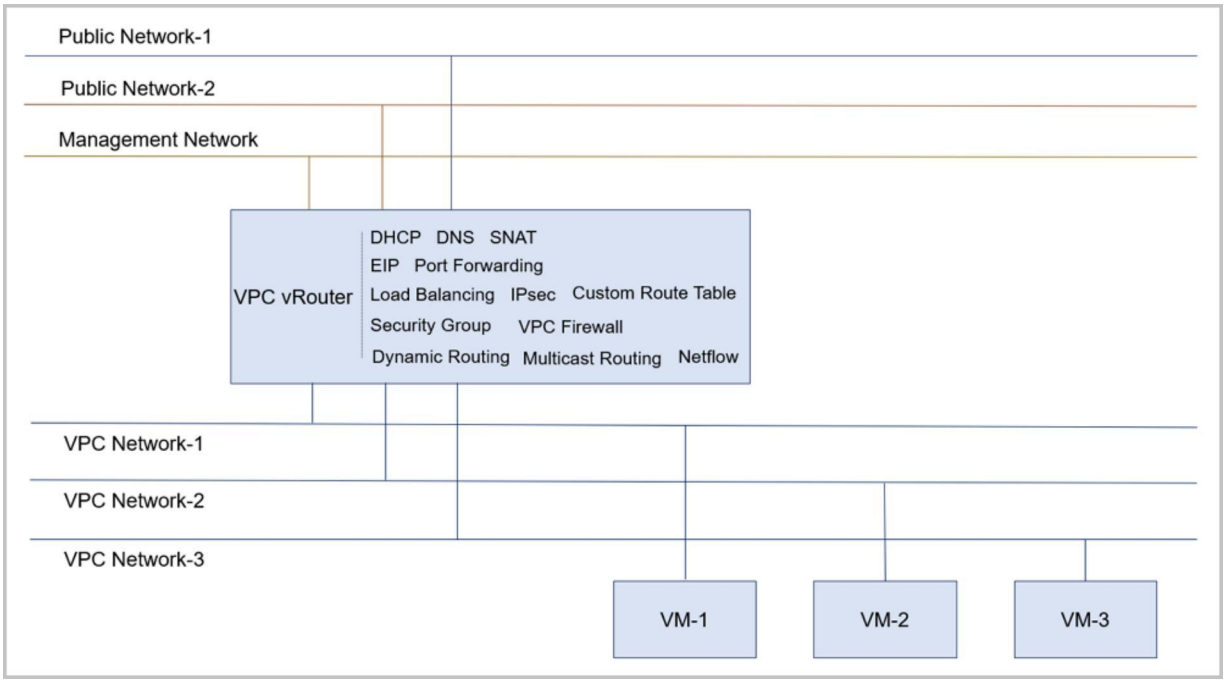

VPC Network Services

The VPC network, which acts as a private network, provides a group of network services by using VPC vRouters.

- DHCP: By default, the VPC network provides distributed DHCP services by using the flat network service module.
- DNS: A VPC vRouter can act as a DNS server to provide DNS services. The DNS address in a VM instance is the IP address of the VPC vRouter. Note that the DNS address that you set is forwarded by the VPC vRouter.
- SNAT: A VPC vRouter can provide the source network address translation (SNAT) services for VM instances. Then, the VM instances can directly access the Internet by using SNAT.
- Route table: You can manage and customize routes through route tables.
- Security group: The security group service is provided by the security group network service module. You can configure and manage firewalls for VM instances by using iptables.
- Elastic IP address (EIP): You can bind an EIP to a VPC network. Then, the public network can interconnect with the private network of the VM instance.
- Port forwarding: The port forwarding service allows a public IP address to interconnect with the private IP address of a VM instance. To be more specific, you can create port forwarding rules to allow external networks to reach specific ports of your VM instances.

- Load balancing: The load balancing service distributes your inbound traffics from a public IP address to a group of backend VM instances. Then, this service automatically checks and isolates the VM instances that are unavailable.
- IPsec tunnel: The IPsec tunnel can be used to achieve interconnection between different virtual private networks (VPNs).
- Dynamic routing: The VPC vRouter supports the Open Shortest Path First (OSPF) routing protocol, which is used to distribute routing information within a single autonomous system.
- Multicast routing: The VPC vRouter forwards the multicast information sent by the multicast source to VM instances, achieving one-to-multi-point communication in the transmission side and receiving side.
- VPC firewall: The VPC firewall filters the south-north traffic on the VPC vRouter ports, effectivel y protecting the VPC communication security and VPC vRouter security.
- Netflow: The Netflow service monitors and analyzes the inbound and outbound traffics of the VPC vRouter NICs. Currently, the following two types of data-flow output format are supported: Netflow V5 and Netflow V9.

Characteristics

A VPC has the following characteristics:

- Flexible network configuration: Different VPC networks can be flexibly attached to the VPC vRouters. You can customize an independent IP range and an independent gateway for each VPC network. VPC vRouters allow you to attach or detach gateways, and also to dynamically configure your route tables and route entries.
- Secure and reliable isolation: Different VPC networks in different VPCs are logically isolated. That is, the VPC networks support VLAN and VXLAN for logical layer 2 isolation, and different VPCs of different accounts will not affect each other.
- Multi-subnet interconnection: Multiple VPC networks under the same VPC can communicate privately and securely with one another.
- Network traffic optimization: VPC supports distributed route features, which can optimize the east-west network traffic and reduce the network latency effectively.
- VPC vRouter HA: In a VPC vRouter HA group, you can deploy two VPC vRouters according to the active-standby policy. When the active VPC vRouter is abnormal, the standby VPC vRouter will automatically take over to work properly, thus ensuring your business continuity.

Routing Protocol Resource

Compared to static routing, dynamic routing, which can be applied to a large-scale network environment, supports automatic topology change, route recalculation, and unattended interferen ce. A VPC vRouter supports the OSPF dynamic routing protocol.

Open Shortest Path First (OSPF): An OSPF is an interior gateway protocol of link states and is used to distribute routing information within a single autonomous system (AS). An OSPF is widely used in a data center network and a campus network.

2.2.1.4 Network Service

The Cloud provides VM instances with multiple network services, including VPC firewall, security group, virtual IP address (VIP), elastic IP address (EIP), port forwarding, IPsec tunnel, load balancing, and flow monitoring.

The Cloud supports the following two network models:

- Flat network
- VPC network

Network Service Module

The Network Service Module provides a group of network services. Note that this module has been hidden on the UI.

The Network Service Module has the following four types:

1. Virtual Router Network Service Module (Not recommended)

   Provides various network services: DNS, SNAT, load balancing, port forwarding, EIP, and DHCP.

2. Flat Network Service Module (Flat Network Service Provider)
   Provides the following network services:

   - User Data: Customizes some parameters, such as `ssh-key` injection. By running `cloud -init`, these parameters will be loaded and injected into your VM instance when the VM instance is started.
   - EIP: Allows you to access private networks through public networks.
   - DHCP: Dynamically obtains an IP address.

Note:

The DHCP service includes the DNS feature.

- VIP QoS: Limits the upstream and downstream bandwidth. This applies only to EIPs.

3. VPC vRouter Network Service Module

   Provides the following network services:

   - IPsec: Achieves VPN connections.
   - vRouter route table: Manages custom routes.
   - Centralized DNS: Provides the DNS service when the distributed DHCP service is enabled.
   - VIP QoS: Limits the upstream and downstream bandwidth of a virtual IP address.
   - DNS: Uses VPC vRouters to provide the DNS service.
   - SNAT: Enables VM instances to access the Internet directly.
   - Load balancing: Distributes inbound traffics from a VIP to a group of backend VM instances . Then, unavailable VM instances will be detected and isolated automatically.
   - Port forwarding: Forwards port traffics of specified public IP addresses to the ports of corresponding VM instances according to specified protocols.
   - EIP: Uses VPC vRouters to access private networks of VM instances through public networks.
   - DHCP: Provides the centralized DHCP service.

4. Security Group Network Service Module
   Provides the following network service:

   - Security group: Manipulates securities of VM instance firewalls by using iptables.

Flat Network Practice

In your production environments, we recommend that you use the following combination of network services:

- Flat Network Service Module
  - User Data: Customizes some parameters, such as `ssh-key` injection. By running `cloud-init`, these parameters will be loaded and injected into your VM instance when the VM instance is started.
  - EIP: Allows you to access private networks through public networks.
  - DHCP: Dynamically obtains an IP address.

Note:

The DHCP service includes the DNS feature.

- Security Group Network Service Module
  - Security group: Manipulates securities of VM instance firewalls by using iptables.

VPC Network Practice

In your production environments, we recommend that you use the following combination of network services:

- Flat Network Service Module
  - User Data: Customizes some parameters, such as `ssh-key` injection. By running `cloud-init`, these parameters will be loaded and injected into your VM instance when the VM instance is started.
  - DHCP: Dynamically obtains an IP address.
- vRouter Network Service Module
  - DNS: Uses vRouters to provide the DNS service.
  - SNAT: Allows VM instances to access directly the Internet.
  - vRouter route table: Manages custom routes.
  - EIP: Uses vRouters to access private networks of VM instances through public networks.
  - Port forwarding: Forwards port traffics of specified public IP addresses to the ports of corresponding VM instances according to specified protocols.
  - Load balancing: Distributes inbound traffics from a VIP to a set of backend VM instances. Then, unavailable VM instances will be detected and isolated automatically.
  - IPsec: Achieves VPN connections.
- Security Group Network Service Module
  - Security group: Manipulates securities of VM instance firewalls by using iptables.

Advanced Network Services

- Dynamic routing: Uses the Open Shortest Path First (OSPF) routing protocol to distribute routing information within a single autonomous system. This service applies to VPC network scenarios.
- Multicast routing: Forwards the multicast information sent by the multicast source to VM instances, achieving one-to-multi-point communication in the transmission side and receiving side. This service applies to VPC network scenarios.

- VPC firewall: Filters the south-north traffic on the VPC vRouter ports, effectively protecting the VPC communication security and VPC vRouter security. This service applies to VPC network scenarios.
- Port mirroring: Copies and sends network traffics of VM NICs from a port to another port, and analyzes the business packets on the ports, better monitoring and managing the network data. This service applies to flat network, vRouter network, and VPC network scenarios.
- Netflow: Monitors and analyzes the inbound and outbound traffics of the VPC vRouter NICs. Currently, the following two types of data-flow output formats are supported: Netflow V5 and Netflow V9. This service applies to VPC network scenarios.

2.2.1.4.1 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.

Characteristics

Security Group and Security Rule

A security group relies on security rules to filter flows accessing or out of VM NICs. You can add one or more security rules to a security group.

- Security rules filter flows based on the flow source or flow destination. They can be categorized into the following two types based on the direction of flows they control:
  - Ingress Rule: Ingress rules take effect on flows accessing VM NICs. They are responsible for filtering ingress flow sources.
  - Egress Rule: Egress rules take effect on flows out of VM NICs. They are responsible for filtering egress flow destinations.
- You can set IP addresses or other security group as flow sources/destinations of security rules.
  - IP address as source: A source IP is filtered by ingress rules. The rules may allow or reject the flows from this IP address to access VM NICs.
  - Security group as source: A source security group is filtered by ingress rules. The rules may allow or reject the flows from this security group to access VM NICs.
  - IP address as destination: A destination IP is filtered by egress rules. The rules may allow or reject VM NICs to access this IP address.
  - Security group as destination：A destination security group is filtered by egress rules. The rules may allow or reject VM NICs to access this security group.

- You can set priorities for rules on the same direction. The highest rule take effect when a conflict occurs in such a scenario as you set more than one rule, especially an allow rule and a reject rule, on the same source or destination.
- By default, mutual communications among NICs in the same security group are allowed and the system automatically add corresponding ingress/egress rules to the security group to ensure these mutual communication. These default rules cannot be modified or deleted. If you want to cancel the mutual communications, just disable these rules.

Security Group and VM NIC

A security group provide security controls to VM NICs attached to it. A security group can be attached to one or more VM NIC, and a VM NIC can be attached to one or more security group.

- If you attach more than one security groups to a VM NIC, you can set priorities for these groups. The NIC matches the rules of the group with the highest priority first, and then the group of lower priorities.

Note:
By default, all admin security group have higher priority than user security groups.

- After attached to security groups, you need to set a default flow policy to process the flows that are not stipulated by security group rules. By default, all ingress rules that are not stipulated are rejected and all egress rules that are not stipulated are allowed.

Security Group and Permission

Security groups are divided into admin security groups and tenant/sub-account security groups. Generally, admin security groups are created and owned by administrators (including admin and platform managers); tenant/sub-account security groups are created and owned by tenants/sub-accounts.

- A tenant/sub-account can view and manage security groups owned by itself.
- The administrator can view and manage all security group. When attach security groups to NICs, note that an admin security group can be attached to any NIC, while To a tenant/sub-account security group can be attached to only NICs owned by the same tenant/sub-account.

Considerations

- If you use a security group along with other network services, such as load balancing and route table, make sure that the security group rules required by these network services are added to the security group.

- Public networks, flat networks, and VPC networks support the security group service. It is provided by the security group network service module, which uses iptables to implement security control.
- A security group is a distributed firewall. Each security rule change, NIC association or disassociation will cause the security group rule to be updated on all associated VM instances.

2.2.1.4.2 Virtual IP

In bridged network environments, a virtual IP address (VIP) provides network services such as serving as an elastic IP address (EIP), port forwarding, load balancing, IPsec tunneling. When a VIP provides the preceding network services, packets are sent to the VIP and then routed to the destination network where VM instances are located.

- VIPs created from a public network can provide network services such as EIP and load balancing for flat networks.
- VIPs created from a public network can provide network services, such as EIP, port forwarding , load balancing, and IPsec tunnel, for VPC networks.
- VIPs created from a VPC network can provide load balancing services for VPC networks.
- VIPs created from a flat network can provide network services, such as EIP and load balancing , for flat networks.

You can use a VIP to provide performance-shared load balancing services. A performance-shared load balancer uses a VPC vRouter to provide load balancing services. Traffic is distributed to backend servers by the VPC vRouter. If the VPC vRouter is providing multiple services, the load balancing service shares the performance of the VPC vRouter along with other services.

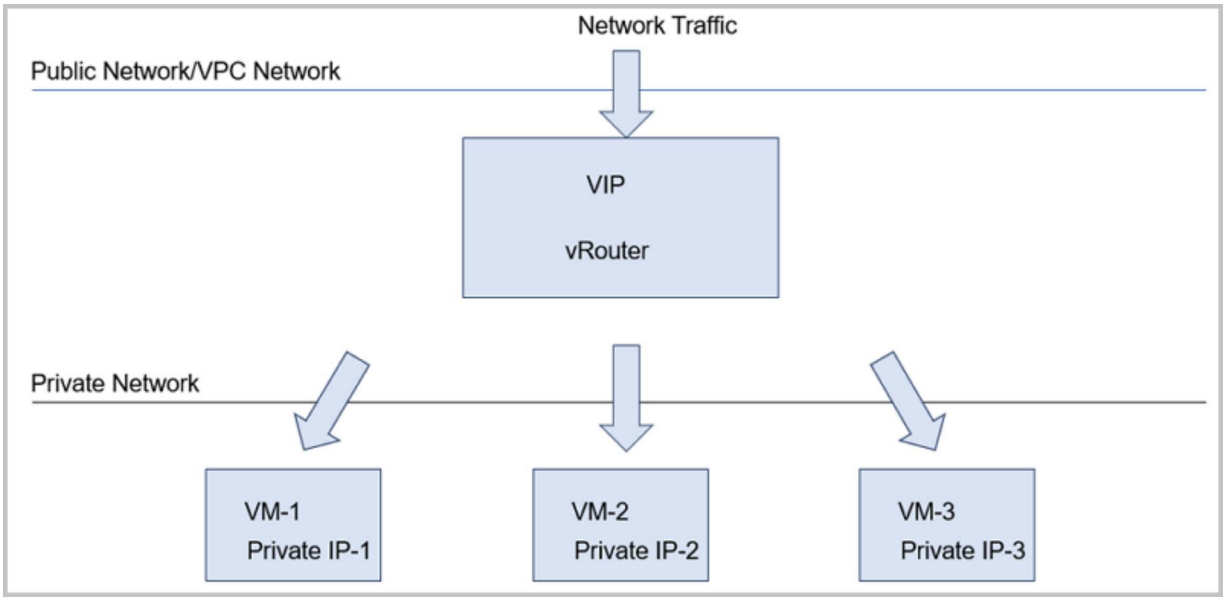

Concepts

- Public VIP: VIPs created from a public network. You can manually create a public VIP or use a public VIP automatically created after a VPC vRouter creation.
  - A public VIP can provide network services, such as EIP and load balancing, for flat networks . A public VIP can also provide network services, such as EIP, port forwarding, load balancing, and IPsec tunnel, for VPC networks.
  - You can use a public VIP to simultaneously provide port forwarding, load balancing, and IPsec tunnel services. You can also use a public VIP to provide one service for multiple VM instances. However, you cannot specify the same port number for different services.
  - A public VIP supports QoS, monitoring data, performance TOP 5, performance analysis, alarm, and other features.
- VPC VIP: VIPs created from a VPC network. A VPC VIP can only be manually created.
  - A VPC VIP can provide load balancing services for VPC networks.
  - VPC VIPs do not support QoS, monitoring data, performance TOP 5, performance analysis, and alarm features.
- Flat network VIP: VIPs created from a flat network. You can manually create a flat network VIP or use a flat network VIP automatically created after a VPC vRouter creation.
  - A flat network VIP provides network services, such as EIP and load balancing, for flat networks.
  - A flat network VIP supports QoS, monitoring data, performance TOP 5, performance analysis, alarm, and other features.
- Custom VIP: manually created VIPs. You can customize a public VIP, VPC VIP, and flat network VIP based on your needs.
  - One custom public VIP can only be applied to one EIP instance.
  - Custom VIPs cannot be used across VPC vRouters.
  - When you use the EIP, port forwarding, load balancing, or IPsec tunnel services, you can select **Create VIP** to create a new VIP, or you can select **Use Existing VIP** to provide the services.
- System VIP: VIPs automatically created by using the L3 network attached to VPC vRouters after the VPC vRouter creation. System VIPS can be categorized into public VIPs and flat network VIPs.

- A system VIP belongs to only one VPC vRouter. When you attach a public network to a VPC vRouter, the Cloud will automatically create a system VIP. This VIP is the default IP address of the vRouter in the network.
- By default, the system VIPs created from public networks are used to provide the source network address translation (SNAT) service.
- When you use the EIP, port forwarding, load balancing, or IPsec tunnel service, you can select **Use Existing VIP** to provide the services.

2.2.1.4.3 Elastic IP

An elastic IP address (EIP) functions based on the NAT technology. IP addresses in a private network are translated into an EIP that is in another network. This way, private networks can be accessed from other networks by using EIPs.

Concepts

- Public EIP: The EIP service provided by a virtual IP address (VIP) created from a public network.
  - The private network is an isolated network that cannot be directly accessed from other networks or the Internet. A public EIP can directly associate the access to a public network with the VM IP of the private network.
  - A public EIP can be associated with or disassociated from a VM instance dynamically.
  - A public EIP can be associated with VM instances created from private networks, such as flat networks and VPC networks.
    - The public EIP realized by a distributed EIP can access flat networks through public networks.
    - A VPC vRouter can be used to access VPC networks through public networks.
- Flat EIP: The EIP service provided by a VIP created from a flat network.
  - L3 isolations exist between flat networks of different network ranges. Therefore, these flat networks cannot be accessed directly. A flat EIP can be used to associate the access to one flat network with the VM IP created from another flat network.
  - A flat EIP can be associated with or disassociated from a VM instance dynamically.
  - A flat EIP can be associated with VM instances created from other flat networks.

Scenarios

- EIP usage in a flat network scenario:

The following diagram shows how an EIP is used in a flat network scenario.

Figure 2-14: EIP Usage in Flat Network Scenario

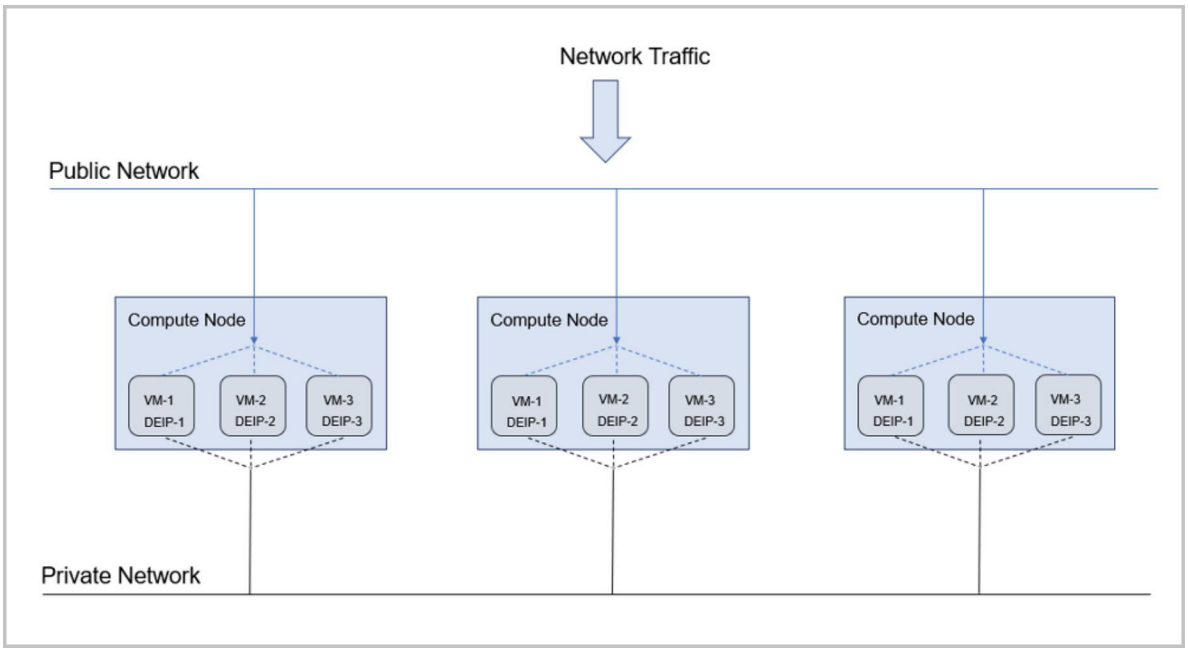

- Public networks can connect to the Internet through firewalls.
- Flat networks provide IP addresses for VM instances in each compute node. Notice that these IP addresses cannot connect to the Internet by default.
- A distributed EIP is deployed on each compute node, and can be associated with public networks or private networks separately.

- EIP usage in a VPC network scenario:

  The following diagram shows how an EIP is used in a VPC network scenario.

Figure 2-15: EIP Usage in VPC Network Scenario

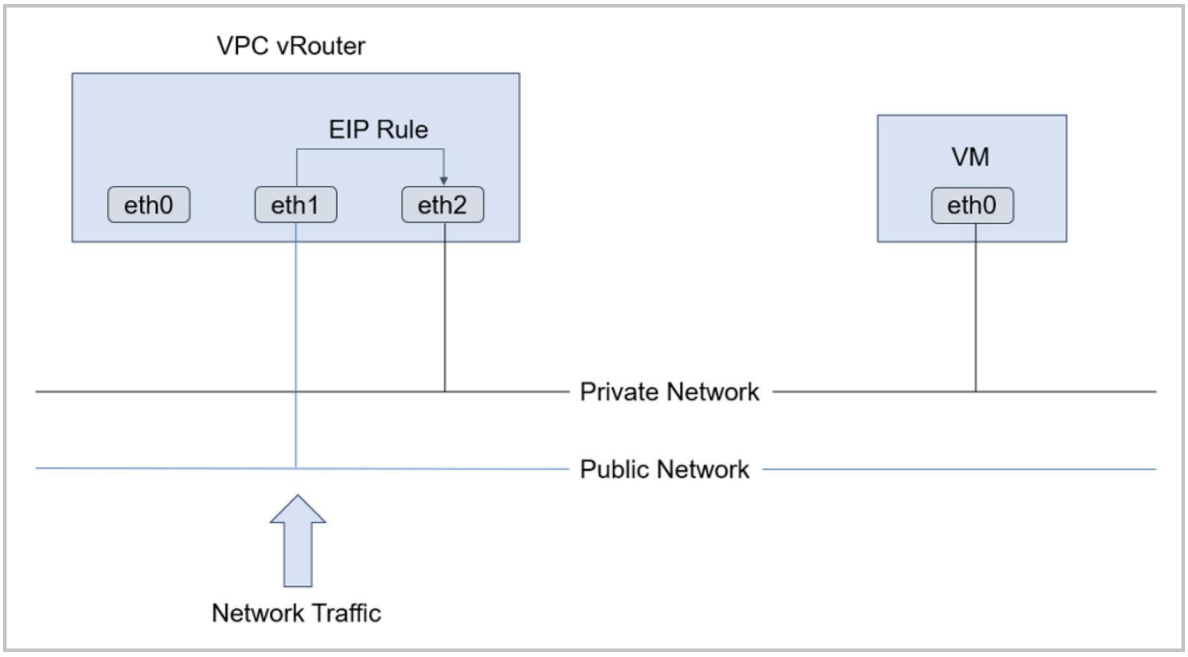

Considerations

When you use an EIP, note that:

- An instance can have only one EIP associated at a time.
- The EIP association and disassociation operations take effect in real time.
- Associating or disassociating an EIP does not affect the running of an instance.

2.2.1.4.4 Port Forwarding

Port forwarding functions based on the layer-3 forwarding service of VPC vRouters. This service forwards traffic flows of the specified IP addresses and ports in a public network to specified ports of VM instances by using the specified protocol. If your public IP addresses are insufficient, you can configure port forwarding for multiple VM instances by using one public IP address and port.

- VM instances in a private network for which SNAT is enabled can access external networks . However, the VM instances in a private network are inaccessible to external networks. You can use a port forwarding rule to allow the access to specified ports of the VM instances from external networks.
- You can associate a port forwarding rule with a VM NIC and disassociate a port forwarding rule from a VM NIC based on your business needs.
- Port forwarding services are provided only by VPC vRouters.

Port forwarding rules are applied to a public network associated with a VPC vRouter and a private network where VM instances reside, as shown in the following figure.

Figure 2-16: Port Forwarding

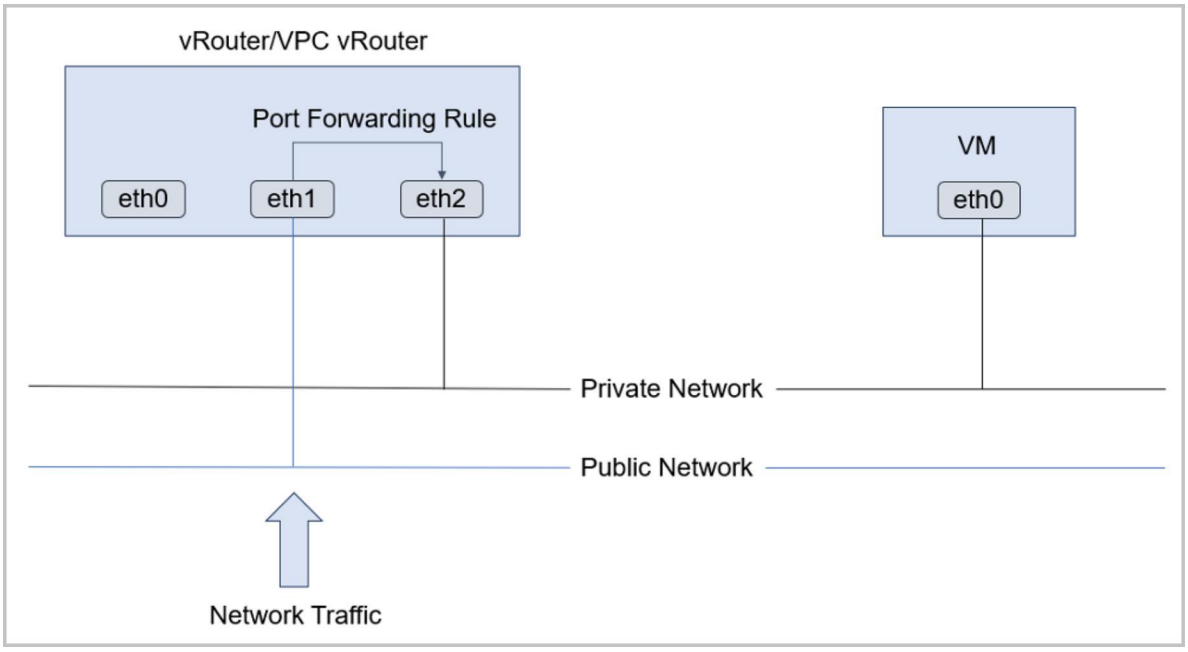

- Port forwarding is achieved by using a virtual IP address (VIP).
  - A VIP is an available IP address in a public network.
  - You can use an existing VIP or create a VIP to provide port forwarding services.
  - Two port forwarding methods are supported: port-to-port mapping and port range-based mapping.

Figure 2-17: VIP-Port Forwarding

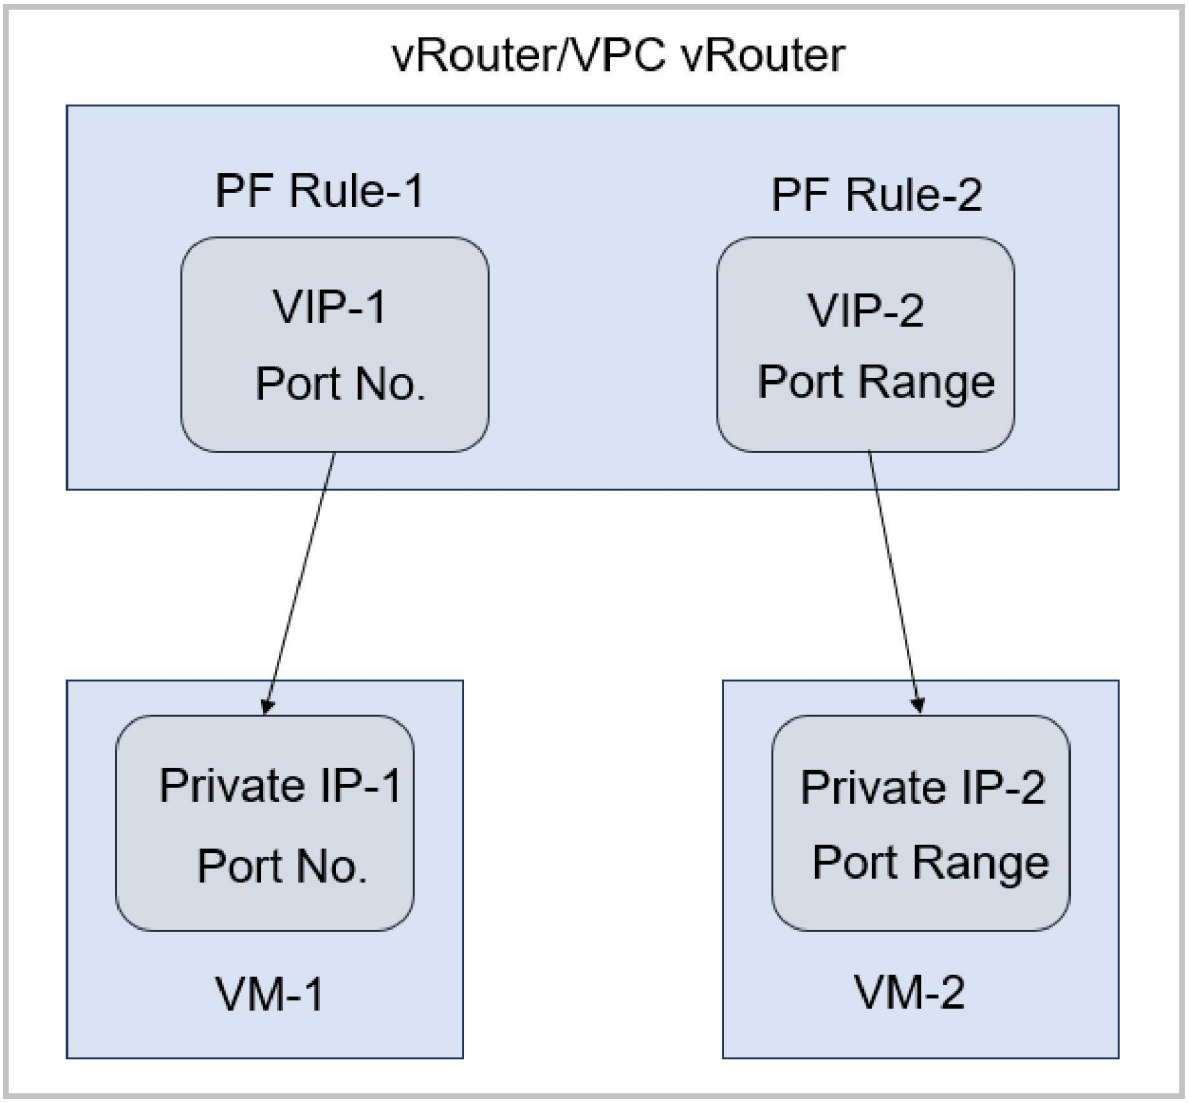

Limits

The port forwarding service has the following limits:

- The firewall policy of a VM instance must open the port specified for port forwarding.
- The ports used for port forwarding by the same VIP must be unique.
- You can use a VIP to provide port forwarding services for different ports of multiple VM NICs in the same L3 network.
- You can use only one VIP to provide port forwarding services for a VM instance.
- If you disassociate a VIP from a VM instance and then associate the VIP with the VM instance, you can select VM NICs that reside in the same L3 network as the previously disassociated VM instance.

- The source port range and target port range used for port forwarding must be consistent. For example, if you set the source port range to 22-80, the target port range must also be 22-80.

2.2.1.4.5 Load Balancing

A load balancer distributes traffic flows of a virtual IP address to backend servers. It automatically inspects the availability of backend servers and isolates unavailable servers during traffic distributi on. This way, the load balancer improves the availability and service capability of your business.

NexaVM Cloud supports multiple NIC teaming for NIC redundancy and load-balancing capabilities. The Cloud provides the following two types of load balancing services:

- Shared-performance load balancing: uses a VPC vRouter to provide load balancing services. Traffic is distributed to backend servers by the VPC vRouter. If the VPC vRouter is providing multiple services, the load balancing service shares the performance of the VPC vRouter with other services.
- Dedicated-performance load balancing: uses a load balancer instance to provide load balancing services. Traffic is distributed to backend servers by the load balancer instance. A load balancer instance is a custom VM instance dedicated to providing load balancing services.

Concepts

- Frontend network: A frontend network is a type of network that is associated with a load balancer. Requests from the network are distributed by the load balancer to backend servers based on a specified policy.
  - Shared-performance load balancer: You can specify a public network or VPC network as a frontend network.
  - Dedicated-performance load balancer: You can specify a public network, flat network, or VPC network as a frontend network.
- Backend network:A backend network is a type of network that is associated with a load balancer. Requests from frontend networks are distributed by the load balancer to servers in the backend network.
  - Shared-Performance Load Balancer:
    - If you specify a public network as the frontend network, you can specify any one of the VPC networks of the VPC vRouter to which the public network is attached as the backend network.

  - If you specify a VPC network as the frontend network, you can specify any one of the VPC networks of the VPC vRouter to which the VPC network is attached as the backend network.

- Dedicated-Performance Load Balancer:
  - If you specify a public network as the frontend network, you can specify a flat network or any one of the VPC networks of the VPC vRouter to which the public network is attached as the backend network.
  - If you specify a flat network as the frontend network, you must specify the same network as the backend network.
  - If you specify a VPC network of a VPC vRouter as the frontend network, you must specify the same network as the backend network. If you need to associate the load balancer with more backend networks, you can attach more NICs to the load balancer instance. Note that the backend networks that you associate are among the other VPC networks of the VPC vRouter.

- Load balancer instance: A load balancer instance is a custom VM instance used to provide load balancing services.
  - The network where the default NIC of a load balancer instance resides is the frontend network of the load balancer instance. The default NIC cannot be detached from a load balancer instance.
  - The networks of the NICs of a load balancer instance, except the default NIC, are the backend networks of the load balancer.
  - The management NIC (if any) cannot be detached.
- LB image: A dedicated-performance load balancer (LB) image encapsulates dedicated-performance load-balancing services and can be used to create load balancer instances . However, a dedicated-performance load balancer image cannot be used to create VM instances.
  - An LB image is a custom image. You can download the image from the Cloud official website and add the image to the platform.
- Load balancer instance: A load balancer (LB) instance offering defines the CPU, memory , image, and management network configuration settings used to create LB instances. LB instances provide load balancing services for the public network, flat network, and VPC network.

- Listener: A listener monitors the frontend requests of a load balancer and distributes the requests to a backend server based on the specified policy. In addition, the listener performs health checks on backend servers.
  - Listeners support the TCP, HTTP, HTTPS, and UDP protocols.
  - A load balancer can be associated with multiple listeners while a listener can be associated with only one load balancer.
  - If the listener uses the weighted round-robin load-balancing algorithm, you can set the weight value for each individual backend server on the backend server group details page.
- Forwarding rule: A forwarding rule forwards the requests from different domain names or URLs to different backend server groups.
  - A forwarding rule is composed of a domain name and URL.
  - A listener can have up to 40 forwarding rules.
  - You can configure a domain-based or URL-based forwarding rule for a load balancer. You can add multiple forwarding rules to a listener and associate these rules with different backend server groups.
  - A forwarding rule is matched by using the exact match and fuzzy match mechanisms. If multiple forwarding rules are matched, the forwarding rule matched through the exact match mechanism is used.
- Certificate: If you select HTTPS for a listener, associate it with a certificate to make the listener take effect. You can upload either a certificate or certificate chain.
- Backend server group: A backend server group is a group of backend servers that handles requests distributed by load balancers. It is the basic unit for traffic distribution by load balancer instances.
  - A load balancer can be associated with multiple backend server groups, while a backend server group can be associated with only one load balancer.
  - A backend server group can be associated with multiple listeners in the same load balancer.
  - Creating a load balancer will automatically create an empty backend server group.
- Backend server: A backend server handles requests distributed by a load balancer. You can add a VM instance on the Cloud or a server on a third-party cloud as a backend server.

2.2.1.4.5.1 Certificate

The Certificate feature complies with the digital certificate protocol. Trusted certificate authorities (CAs) issue digital certificates after verifying the identity of a server. The issued certificates can verify server identities and encrypt data transmission.

2.2.1.4.6 VPC Firewall

A firewall is an access control policy that monitors ingress and egress traffic of VPC vRouters and decides whether to allow or block specific traffic based on the associated rule sets and rules.

Concepts

- Firewall rule set: A firewall rule set is a set of rules that a firewall uses to defend against network attacks. You need to associate a rule set with the egress or ingress flow direction of VPC vRouter NICs to make the rule set take effect.

  — You can associate a rule set with the egress or ingress flow direction of VPC vRouter NICs:

    - Ingress: applies to the traffic that flows into the specified VPC vRouter via a network.
    - Egress: applies to the traffic that flows out of the specified VPC vRouter via a network.

- Firewall rule: A firewall rule is an access control entry associated with the egress or ingress flow direction of VPC vRouter NICs to defend against network attacks. A firewall rule includes rule priority, match condition, and behavior.

  — You can associate a rule with the egress or ingress flow direction of VPC vRouter NICs:

    - Ingress: applies to the traffic that flows into the specified VPC vRouter via a network.
    - Egress: applies to the traffic that flows out of the specified VPC vRouter via a network.

  — Firewall rules can be categorized into custom rules and system rules:

    - Custom rules: rules that you customize. You can select the ingress or egress direction that the rules take effect and configure the rule priorities, match conditions, and behaviors.

      - Rule priority: the priority of a rule to be matched and take effect when compared with other firewall rules. Valid values: 1001 to 2999.

        - Generally, a rule with a higher priority is primarily matched when compared to a rule with a lower priority. Priorities are represented by using numbers. A smaller number indicates a higher priority.

- Generally, the more specific the match condition that you configure for a rule is, the higher priority you shall configure for the rule.

- Match condition: the condition based on which traffic flowing into or out of a VPC network is matched. It includes source IP address, destination IP address, source port, destination port, packet status, and protocol.

  - You can specify one or more source and destination IP addresses. These IP addresses can be static IP addresses, IP ranges, CIDR blocks, or a mix of the three.
  - If you specify multiple entries, which include one or more CIDR blocks, the netmask of the CIDR block must be 24. If you specify only one CIDR block, the netmask of the CIDR block is not limited.
  - You can enter a maximum of ten entries, with each entry separated by a comma (,).

- Behavior: the action to be applied to traffic that meets the match condition. Valid values: accept, drop, and reject.

  - Accept: accepts the traffic that flows in or out of the specified VPC vRouter.
  - Drop: drops the traffic that flows in or out of the specified VPC vRouter and does not respond to the client.
  - Reject: rejects the traffic that flows in or out of the specified VPC vRouter and responds to the client.

- System rules: rules predefined to support system services. The system predefines the direction that the rules take effect, and the priority, match condition, and behavior of the rules.

  - The priority of system rules ranges from 1 to 1000 or from 4000 to 9999.
  - NexaVM Cloud has predefined the following system rules:

    - Firewall rules that take effect on the ingress direction of VPC vRouter NICs:

      - Rule 1: The priority is 4000, and the behavior and match condition combination determines to allow **established** or **related** data packets from any IP address/port, with any protocol, or to any IP address/port, to flow into the specified VPC vRouter via a network.
      - Rule 2: The priority is 9999, and the behavior and match condition combination determines to allow **new** data packets from any IP address/port, with any

protocol, or to any IP address/port, to flow into the specified VPC vRouter via a network.

- Rule 3: the default rule with a priority of 10000. The behavior and match condition combination determines to reject data packets from any IP address/port, with any protocol, in any status, or to any IP address/port, from flowing into the specified VPC vRouter via a network. You can modify the behavior of the rule. Valid values: accept, drop, and reject.

- Firewall rules that take effect on the egress direction of VPC vRouter NICs:
  - Rule 1: the default rule with a priority of 10000. The behavior and match condition combination determines to reject data packets from any IP address/port, with any protocol, in any status, or to any IP address/port, from flowing into the specified VPC vRouter via a network. You can modify the behavior of the rule. Valid values: accept, drop, and reject.
- System rules cannot be modified, except the behavior of the default rule.
- System rules cannot be created or deleted.

- Rule template: A rule template is a template that you can select when you add rules to a rule set or a firewall.
- IP/Port set: An IP or port set is a set of IP addresses or ports that you can select when you add rules to a rule set or a firewall.

Fundamentals

NexaVM Cloud allows you to associate rule sets and rules with the ingress and egress direction of VPC vRouter NICs. Then traffics that flow in or out of the VPC vRouter NICs are filtered based on the rule priority, match condition, behavior, and the effect direction. This ensures the security of data communications across VPC networks, of VPC vRouters, and of user business operations.

Figure 2-18: Firewall

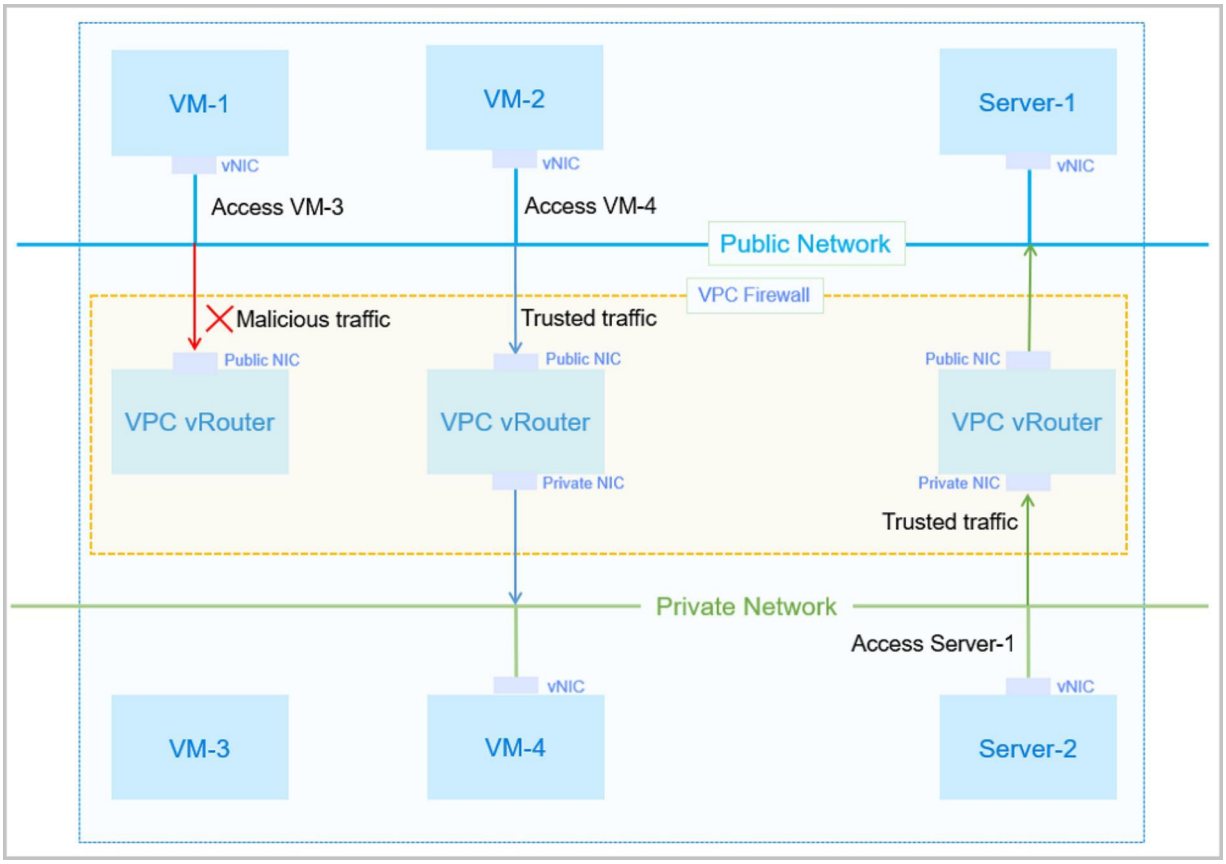

Assume that a server and two VM instances are deployed in a VPC network to run significant business applications. To ensure business security, firewall rule sets and rules are associated with the ingress or egress direction of VPC vRouters, so that only trustful traffics from the public network are allowed to access VM data in the VPC network and that the server in the VPC network can access the server data in the public network.

- When VM-1 attempts to access VM-3: The traffic from VM-1 will match the inbound rule set of the public NIC on the VPC vRouter. If malicious traffics are detected, the access is denied.
- When VM-2 attempts to access VM-4: The traffic from VM-2 will match the inbound rule of the public NIC on the VPC vRouter, and then will match the outbound rule set of the private NIC on the VPC vRouter. If trusted traffics are detected, the access is allowed.
- When Server-2 attempts to access Server-1: The traffic from Sever-2 will match the inbound rule set of the private NIC on the VPC vRouter, and then will match the outbound rule set of the public NIC on the VPC vRouter. If trusted traffics are detected, the access is allowed.

Firewall vs Security

A firewall manages the south-north traffic of VPC networks. A security group manages the east-west traffic of VPC networks and is applied to VM NICs. The two services complement with each other. The following table compares the two services from three aspects.

Item	Security Group	Firewall
Application scope	VM NIC	The entire VPC network
Deployment mode	Distributed	Centralized
Deployment location	VM instance	VPC vRouter
Configuration policy	Supports only Allow policies	Allows you to customize Accept, Drop, or Reject policies as needed
Priority	Takes effect based on the predefined rule sequence	Allows you to customize priorities
Match condition	Source IP address, source port, and protocol	Source IP address, source port, destination IP address, destination port, protocol, and packet status

2.2.1.4.7 IPsec Tunnel

An IPSec tunnel encrypts and verifies IP packets that transmit over a virtual private network (VPN) from one site to another.

The following are the characteristics of an IPsec tunnel:

- IPsec negotiation mode:

  For security reasons, we only support the Main mode. The Aggressive mode is not supported.
- IPsec security protocol:

  We support only the Encapsulating Security Payload (ESP) protocol.
- IPsec encapsulation mode:

  We support the Tunnel mode. The Transport mode is not supported.
- IPsec routing model:

  We support only policy-based IPSec VPN. Route-based IPSec VPN is not supported.

  Therefore, the tunnel supports only unicast data, and does not support multicast and broadcast .

The typical usage scenario of an IPsec tunnel in vRouter networks is as follows:

- Prepare two isolated NexaVM Cloud and set up two VPC environments in these two clouds respectively. In each VPC environment, create two VPC networks respectively and make sure that these VPC networks cannot communicate with each other. Then, you can use an IPsec tunnel to achieve communication between these VPC networks.

2.2.1.4.8 Netflow

A NetFlow monitors the ingress and egress traffic of the NICs of VPC vRouters. The supported versions of data flows are V5 and V9.

2.2.1.4.9 Port Mirroring

Port mirroring mirrors the traffic data of VM NICs and sends the traffic data to the target ports. This allows for the analysis of data packets of ports and simplifies the monitoring and management of data traffic and makes it easier to locate network errors and exceptions.

2.2.1.5 CloudFormation

CloudFormation is a service that simplifies the management of cloud resources and automates deployment and O&S. You can create a stack template to configure cloud resources and their dependencies. This way, resources can be automatically configured and deployed in batches.

CloudFormation provides easy management of the lifecycle of cloud resources and integrates automatic O&S into API and SDK.

Figure 2-19: CloudFormation

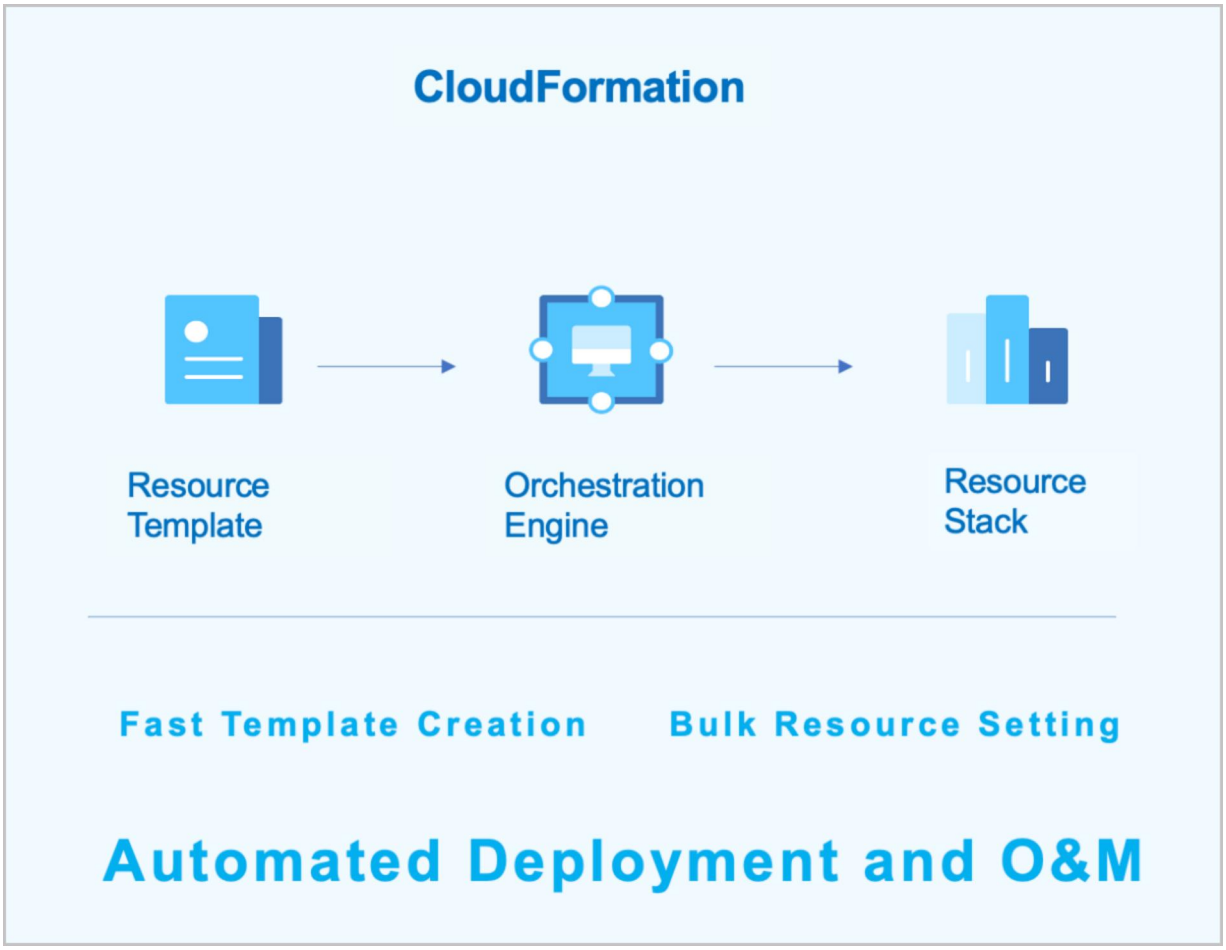

Characteristics

- You can create a stack template or modify an existing one to define what cloud resources you need, the dependencies between the resources, and the resource configurations. Then CloudFormation automatically uses the orchestration engine to create and configure the resources.
- You can use sample templates and the designer provided by the Cloud to create stack templates. This greatly improves efficiency.
- You can update a stack template as needed. Then you can use the updated template to adjust your resource stacks to accommodate the dynamic changes of your business needs.
- If you no longer need a resource stack, you can delete the stack and all resources in it with one click.
- You can use an existing stack template to quickly duplicate all stack resources and their configuration settings.

- You can flexibly combine cloud services according to your business scenarios to realize the automatic maintenance.

2.2.1.6 Baremetal Management

NexaVM Cloud offers the Baremetal Management service that provides your applications with dedicated physical servers, ensuring the high performance and stability of your key applications. After your servers are configured well and the related preparations are completed, you can deploy baremetal (BM) chassis in bulk on the UI. After the deployment succeeds, you can use these BM chassis to create BM instances. With preconfigured templates, you can achieve unattended batch installation for BM instance operating systems. In addition, you can configure a business network for BM instances and easily manage the entire lifecycle of these BM instances.

The Baremetal Management service is a separate feature module. To use this service, purchase both the Base License and the Plus License of Baremetal Management. The Plus License cannot be used independently.

Basic Workflow

How does the Baremetal Management service work? A deployment server provides two types of service: DHCP and FTP. Specifically, the deployment server can instruct multiple BM chassis to be started through a PXE NIC, and can allocate dynamic IP addresses with the DHCP service. In addition, BM chassis can download related software packages from the deployment server with the FTP service, of whose packages can be applied to the operating system installation of the BM instance, as shown in *Baremetal Management Network Topology*.

Figure 2-20: Baremetal Management Network Topology

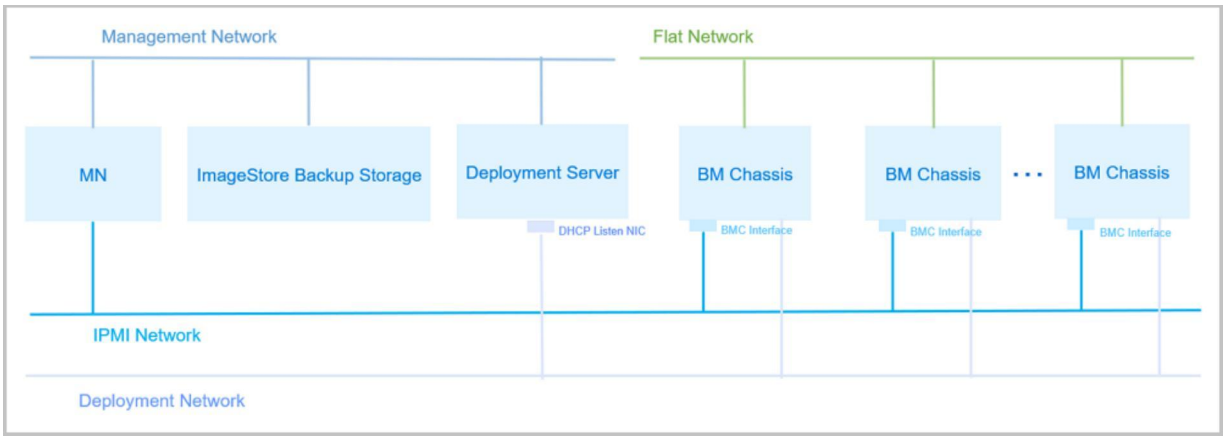

Key Features and Benefits

The Baremetal Management service provides the following features and benefits:

- Provides applications with dedicated physical servers to ensure the high performance and stability of your key applications.
- Deploy deployment servers independently as recommended, which can meet the requirements for the host high availability scenario of multiple management nodes. This also simplifies the network environment and helps to avoid DHCP conflicts. In addition, you can attach an independent deployment server to each BM cluster, which helps to avoid a single point of failure and improve greatly the deployment rate.
- Helps create BM chassis in bulk on the UI via either the manual creation or template file import. You can also add IPMI addresses in bulk to deploy efficiently BM clusters, which increases O&M efficiencies.
- Helps to quickly generate configuration files by using preconfigured templates to achieve unattended batch installation for BM instance operating systems.
- Enables you to customize the installation of your operating system. Supported operating systems: custom operating system of the Cloud and the mainstream Linux distributions (RHEL/CentOS, Debian/Ubuntu, and SUSE/openSUSE).
- Supports flat network. Specifically, BM instances and VM instances on the same L2 network can reach each other without routing to each other by gateways.

Typical Usage Scenarios

The BareMetal Management service can be applied to the following typical scenarios:

- High-Security and Strict Management Scenario

  Financial industry, security industry, and others have rigorous standards for the business compliance and business data security. With the Baremetal Management service, they can ensure their exclusive use of resources, data isolation, strict supervision and control, and effective tracking.

- High-Performance Computing Scenario

  In high-performance computing scenarios, supercomputing centers, gene sequencing companies, and other entities require high computing performance, high stability, and accurate real-time update for servers. Sometime later, business performances will be affected by the performance loss and hyper-threading brought by visualization. In this regard, to deploy BM cluster to a certain scale, meet the strict requirements for the high performance computing.

- Key Database Scenario

  In some entities, some key database businesses cannot be deployed on normal VM instances, and must be loaded on the physical servers that can protect their exclusive resources, network isolation, and performances. To meet this requirements, use the Baremetal Management service that provides exclusive, high-performance physical servers for one or more appliances.

2.2.1.6.1 Baremetal Cluster

A baremetal cluster provides independent cluster managements for baremetal chassis.

- To provide PXE services for baremetal instances on a baremetal cluster, the baremetal cluster must attach a deployment server.
- One baremetal cluster can only attach one deployment server, while one deployment server can attach multiple baremetal clusters simultaneously.
- To provide network services for baremetal instances on a baremetal cluster, the baremetal cluster must attach L2 networks.
- Flat networks are supported. Specifically, both baremetal instances and VM instances on the same L2 network can reach each other without routing via a gateway.

2.2.1.6.2 Deployment Server

A deployment server, known as PXE server, is an independently specified server used for providing PXE services and console proxy services for baremetal chassis.

- We recommend that you deploy PXE servers independently, thus satisfying the need of the multi-MN host HA scenario and avoiding a single point of failure (SPOF) to greatly improve deployment efficiencies.
- A deployment server must be attached to a baremetal cluster.
- One baremetal cluster can only attach one deployment server, while one deployment server can attach multiple baremetal clusters simultaneously.
- A deployment server must have sufficient storage space to save images used for PXE deployments.
- A deployment server must connect to a management network for reaching management nodes.
- A deployment server must connect to a deployment network for reaching baremetal chassis.
- A DHCP listening NIC on a deployment server must connect to a deployment network. In addition, make sure that this deployment network does not contain other DHCP services for avoiding IP conflicts.

- A deployment server must install the latest NexaVM Cloud ISO with the recommended h79c version. Otherwise, this deployment server cannot provide software packages for baremetal chassis via the FTP service.

2.2.1.6.3 Baremetal Chassis

A baremetal chassis can be used to create baremetal instances and can be universally identified via a BMC interface and IPMI configurations. With an IPMI network, a management node can control remotely powers of baremetal chassis, start networks, and enable disks. An administrator can complete deploying all baremetal chassis in bulk on the UI.

- A management node must connect to an IPMI network and control remotely baremetal chassis via IPMI.
- A baremetal chassis must contain a BMC interface, and configures an IPMI address, port, user name, and password, to connect to an IPMI network.
- A deployment server-enabled NIC on a baremetal chassis must connect to a deployment network.
- Other NICs of baremetal chassis can connect to the corresponding L2 networks as needed.

2.2.1.6.4 Preconfigured Template

A preconfigured template can be used to quickly generate a preconfigured file to install baremetal instance operating systems in bulk without attended interferences.

- Make sure that you prepare well the preconfigured template in advance on the Cloud.
- A preconfigured template includes the following two types of templates:
  - System template: Is defaulted by the Cloud, including basic system variables, thereby satisfying a simple, unattended deployment scenario.
  - Custom template: Enable you to upload custom template files with the UTF8 format. Apart from basic system variables, you can customize other variables as needed to satisfy a complex, unattended deployment scenario.

2.2.1.6.5 Baremetal Instance

A baremetal instance is a VM instance created by a baremetal chassis. After you add a baremetal chassis, you can use the baremetal chassis to create baremetal instances.

- A preconfigured template can be used to quickly generate a preconfigured file to install baremetal instance operating systems in bulk without attended interferences.

- Operating system installations can be customized. Currently, the supported versions of operating systems include custom operating systems of the Cloud and mainstream Linux distributions (RHEL/CentOS, Debian/Ubuntu, and SUSE/openSUSE). These versions must be ISO and non-live CD.
- A business network can be configured for baremetal instances. Currently, flat networks are supported. Specifically, both baremetal instances and VM instances on the same L2 network can reach each other without routing via gateways. Make sure that a baremetal cluster where the baremetal chassis resides attaches the corresponding L2 network in advance.

2.2.1.7 Elastic Baremetal Management

Elastic Baremetal Management provides dedicated physical servers for your applications to ensure high performance and stability. In addition, this feature allows elastic scaling. You can apply for and scale resources based on your needs. Elastic Baremetal Management integrates the benefits of hosts and VM instances. It delivers powerful and stable computing capacities of hosts and allows you to use primary storages, L3 networks, and other resources on the Cloud for your applications. This avoids virtualization overheads and improves the availability of cloud resources, allowing you to flexibly use cloud resources as well as physical resources. You can use this feature for application deployment in traditional non-virtualization scenarios.

- The Elastic Baremetal Management feature is provided in a separate module. Before you can use this feature, you need to purchase the Plus License of Elastic Baremetal Management, in addition to the Base License.
- A tenant can use an elastic baremetal offering shared by the admin to create an elastic baremetal instance.

Concepts

- Provision network: A provision network is a dedicated network for PXE boot and image downloads while creating elastic baremetal instances.
  - Before you can use Elastic Baremetal Management, you need to deploy an IPv4 provision network.
  - Provision networks require high network performance. We recommend that you use at least 10 Gigabit NICs for your provision network.
  - You can configure a gateway for your provision network. This way, the provision network can be connected to other networks. If you do not need to connect your provision network to other networks, you do not need to configure a gateway for your provision network.

- Elastic baremetal cluster: An elastic baremetal cluster consists of elastic baremetal instances. You can manage elastic baremetal instances by managing an elastic baremetal cluster where the instances reside.
  - You must attach a provision network to an elastic baremetal cluster to provide PXE services for baremetal nodes in the cluster.
  - You can attach only one provision network to an elastic baremetal cluster. However, you can attach a provision network to multiple elastic baremetal clusters.
  - You can attach an L2 network to an elastic baremetal cluster to provide an extended L2 business network for elastic baremetal instances in the cluster. Elastic baremetal instances and VM instances that share the same L2 network can access each other without using the gateway. The L2 network that you can attach to an elastic baremetal cluster can be of the VLAN or NoVLAN type.
- Gateway node: A gateway node is a node where the ingress and egress traffic of the Cloud and elastic baremetal instances is forwarded.
  - You can attach multiple gateway nodes to an elastic baremetal cluster. However, you can attach only one gateway node to an elastic baremetal cluster.
  - A gateway node is used to take over primary storages and assign storage space for elastic baremetal instances.
  - A gateway node provides iPXE, DHCP, and other services. It is used to deliver configuration settings to elastic baremetal instances.
- Baremetal node: A baremetal node is used to create a baremetal instance and is identified based on the BMC interface and IPMI configuration setting.
  - You can set the startup methods for the baremetal nodes. The following two startup methods are supported:
    - Startup from a volume: Uses a volume as the system volume of the baremetal node to install and deploy the operating system.
    - Startup form a local disk: Uses a local disk as the system volume of the baremetal node to install and deploy the operating system.
    - If you select to startup the baremetal node from a local disk, you can choose whether to take over the original operating system.
      - If you choose not to take over the original operating system, the Local Disk (Non Take -Over) is used. When you use the baremetal node to create an elastic baremetal

instance, this method helps download an operating system from the Cloud and install it on the instance. Meanwhile, the local system volume is formatted.

- If you choose to take over the original operating system, the Local Disk (Take-Over) method is used. The elastic baremetal instance created from the baremetal node used the original system operating system stored on the local disk directly.

- The management node must be connected to the IPMI network to remotely manage baremetal nodes.
- Baremetal nodes must be configured with the BMC interfaces, IPMI addresses, ports, usernames, and passwords, and be connected to the IPMI network.
- A baremetal node can be distributed to only one elastic baremetal instance and an elastic baremetal instance can only be assigned one baremetal node.
- You can provide compute resources for elastic baremetal instances by using a baremetal node or elastic baremetal offering.

- Elastic baremetal instance offering: An elastic baremetal offering defines the number of vCPU cores, memory size, CPU architecture, CPU model, and other configuration settings of elastic baremetal instances.
  - You can get elastic baremetal offerings of baremetal nodes by obtaining their hardware information. Baremetal nodes with the same offering can be managed in a unified way.
  - The baremetal offerings obtained from the node hardware information can be classified into 3 types according to their startup methods: startup from a volume, startup from a local disk (take-over), and startup from a local disk (non take-over).
  - You can use an elastic baremetal offering to create an elastic baremetal instance. You can also release the advanced settings of baremetal nodes to avoid resource idling.
  - You can create a pricing list for elastic baremetal instances based on elastic baremetal offerings. Then bills are generated for the elastic baremetal instances based on their usage.
- Elastic baremetal instance: An elastic baremetal instance has the same performance as physical servers and allows elastic scaling. You can apply for and scale resources based on your needs.
  - The following two startup methods are supported for elastic baremetal instances:
    - Volume: Uses a volume as the system volume of the elastic baremetal instance to install and deploy the operating system.
    - Local Disk: Uses a local disk as the system disk of the elastic baremetal instance to install and deploy the operating system.

- The Local Disk (Non Take-Over) and Local Disk (Take-Over) methods are supported:
  - Local Disk (Non Take-Over): When you use a baremetal node to create the elastic baremetal instance, the operating system is downloaded from the Cloud and installed on the elastic baremetal instance. This method will format the local system disk.
  - Local Disk (Take-Over): When you use a baremetal node to create the elastic baremetal instance, the original operating system on the local system disk is used as the operating system of the elastic baremetal instance.

— The following describes the resources supplied to elastic baremetal instances of different startup methods:

- To elastic baremetal instances of both the volume startup method and local disk startup method, the compute resources are provided by corresponding baremetal nodes, and L3 networks on the Cloud are used as their business networks.
- To elastic baremetal instances of the volume startup method, the storage resources are provided by SharedBlock or Ceph primary storage on the Cloud, and the PXE boots are supported by the provision network.
- To elastic baremetal instances of local disk startup method, the storage resources are provides by their local disks. If you attach data volumes to the instances, they can use the storage resources provided by SharedBlock or Ceph primary storage on the Cloud.

— We recommend that you create an elastic baremetal instance by using an image that has installed the agent. Otherwise, you cannot perform the following actions on the instance: open the instance console, modify the instance password, attach a volume to or detach a volume from the instance, and attach a network to or detach a network from the instance.

— By default, you can you an image whose BIOS mode is UEFI to create an elastic baremetal instance. If you need to use an image with a Legacy BIOS mode, contact the official technical support.

— You can configure business networks for elastic baremetal instances. If you attach an L2 network to the cluster where your baremetal nodes reside, elastic baremetal instances and VM instances that share the same L2 network can access each other without using the gateway.

— You can enable the elastic baremetal instances of the volume startup method to automatically release the associated baremetal node when it is powered-off. The baremetal node released can be used by other elastic baremetal instances, thus avoiding the resource idling .

Scenarios

- Scenarios that require high security and strict monitoring:

  The financial and insurance industries have high requirements over business deployment compliance and data security. In these scenarios, you can use Baremetal 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 scenarios:

  Supercomputing, genome sequencing, and other high-performance computing scenarios have high requirements over the computing performance, stability, and timeliness. However, the virtualization may cause performance losses and hyperthreading may negatively influence the business. Deploying a reasonable number of baremetal clusters can solve these problems, meeting the high-performance computing requirements.

- Key database scenarios:

  To meet business requirements, you may not want to deploy some key databases on VM instances while want to deploy the databases on physical servers that feature dedicated resources, network isolation, and guaranteed performance. In these scenarios, you can use Baremetal Management to provide dedicated high-performance physical servers for your applications.

Advantages

Elastic Baremetal Management has the following advantages:

- Integration of high performance and scalability:

  Elastic Baremetal Management provides dedicated physical servers for your applications to ensure high performance and stability. In addition, this feature allows elastic scaling. You can apply for and scale resources based on your needs.

- Strong scalability:

  A single management node allows you to manage 10,000 baremetal nodes and scale-out according to your needs.

- Strong compatibility:

  An elastic baremetal instance no longer depends on IPMI/BMC network (optional), and the NIC can be used as a baremetal node as long as it supports PXE boot. The realization of the

elastic baremetal instance does not associate any CPU or virtualization technology. As long as there is a system image that can be started, and no matter if it is an ISO, qcow2, or raw image, it can be used as a virtual hard disk by an elastic baremetal instance. That is, the elastic baremetal can be directly used on domestic servers.

In addition, the elastic baremetal management feature can be achieved by software only, without purchasing any proprietary hardware. It is compatible with all x86 and most of domestic ARM CPU architectures and supports mainstream x86 operating systems and some ARM operating systems.

- Advanced technology:

  In addition to turn on/off baremetal servers and deploy system, elastic baremetal management also supports the use of virtual resources on the Cloud, including VPC/flat/public networks, volumes, primary storage and other resources, which seamlessly connects physical resources and cloud resources, and greatly improves the availability of cloud resources.

- Consistency in operation experience:

  All features of a VM instance can be directly operated on an elastic baremetal instance. And an elastic baremetal instance can use any advanced network services without purchasing additional network hardware.

- Flexible deployment:

  An elastic baremetal instance supports both volume deployment method and local disk deployment method, respectively using the primary storage resource and the local disk on the Cloud, which combines the scalability of resources on the Cloud with the stable I/O and high throughput of local disks.

  The primary storage resources used by an elastic baremetal instance can be shared by the KVM clusters of the Cloud, which relieves you from deploying additional storage resources. In addition, local disk deployment method supports take over the original system, which effectively ensures the business continuity.

2.2.1.8 VMware Management

Introduction

VMware Management manipulates VMware vCenter via VMware public APIs and seamlessly integrates some features of VMware vCenter Server to achieve a unified management of multiple virtualization platforms on NexaVM Cloud. If you deployed a VMware vCenter Server, you can

use VMware Management of NexaVM Cloud to take over the VMware vCenter Server. Then you can

view vSphere servers and virtual machines managed by the VMware vCenter Server, use VMware vSphere resources in the virtual data center, and perform operations on virtual machines in the VMware vCenter clusters.

- The VMware Management feature is provided in a separate module. Before you can use this feature, you need to purchase the Plus License of VMware Management, in addition to the Base License.
- NexaVM Cloud supports multiple vCenter versions, including 5.5, 6.0, 6.5, 6.7, and 7.0.

Basic Resource

NexaVM Cloud can manage vCenter basic resources in a unified manner. In NexaVM Cloud, you can add a vCenter, synchronize data for a vCenter, and delete a vCenter.

After you add a vCenter for the first time, NexaVM Cloud will automatically synchronize the clusters, hosts, VM instances, templates, storages, networks, and other resources in the vCenter. To use a managed vCenter, click **Sync Data** to synchronize vCenter resources to your current Cloud. Then, you can view these resources in the UI.

- You can add and manage multiple vCenters.
- You can filter resources before you import vCenter resources to NexaVM Cloud.
  - dvSwitch scenario:

    Only resources of the hosts that are added to a dvSwitch can be imported to NexaVM Cloud. If you do not add a host to a dvSwitch, the associated resources cannot be imported to

    NexaVM Cloud.
  - vSwitch scenario:

    The resources imported to NexaVM Cloud must be the resources of the hosts in the same cluster and must be added with at least one same vSwitch name. In addition, they must have at least one same port group attribute (including the same network labels and the same VLAN ID) at the same time.

Note:

NexaVM Cloud can only manage VM networks rather than VMkernels or management networks.

VM Instance

After you add a vCenter, VM instances in the vCenter will be automatically synchronized to NexaVM Cloud. You can also create vCenter VM instances on your Cloud.

Network

Before you can create new VM instances in the vCenter managed by NexaVM Cloud, create a VPC network or a flat network in the vCenter in advance.

vCenter network service currently supports the VPC network architecture model.

A VPC network provides network services such as SNAT, DHCP, elastic IP (EIP), port forwarding, load balancing, and IPsec tunnel.

- SNAT: A VPC vRouter provides the source network address translation (SNAT) service to vCenter VM instances. vCenter VM instances can directly access the Internet by using SNAT.
- DHCP: Centralized DHCP services realize a dynamic IP address obtainment.
- EIP: Allows a VPC vRouter to access the private network of a vCenter VM instance through a public network.
- Port forwarding: Forwards the port traffics of a specified public IP address to the port of a corresponding vCenter VM IP address.
- Load balancing: Distributes inbound traffics from a public IP address to a group of backend vCenter VM instances, and then automatically detects and isolates unavailable vCenter VM instances.
- IPsec tunnel: Uses an IPsec tunnel protocol to provide site-to-site VPN connections.

Network Service

A VPC network provides network services such as SNAT, DHCP, elastic IP (EIP), port forwarding, load balancing, and IPsec tunnel.

- SNAT: A VPC vRouter provides the source network address translation (SNAT) service to vCenter VM instances. vCenter VM instances can directly access the Internet by using SNAT.
- DHCP: Centralized DHCP services realize a dynamic IP address obtainment.
- EIP: Allows a VPC vRouter to access the private network of a vCenter VM instance through a public network.
- Port forwarding: Forwards the port traffics of a specified public IP address to the port of a corresponding vCenter VM IP address.

- Load balancing: Distributes inbound traffics from a public IP address to a group of backend vCenter VM instances, and then automatically detects and isolates unavailable vCenter VM instances.
- IPsec tunnel: Uses an IPsec tunnel protocol to provide site-to-site VPN connections.

NexaVM Cloud supports multi-account management in a managed vCenter. Normal accounts and project members can use vCenter network services, including EIP, port forwarding, and load

balancing.

Volume

vCenter volumes provide storages for vCenter VM instances. Volumes can be divided into following two types:

- Root volume: A system volume of the VM instance. The root volume supports the system operation of a VM instance.
- Data Volume: A data volume used for the VM instance. The data volume provides extended storage space for a VM instance.

Volume management in vCenter mainly involves data volumes. 。

Image

In NexaVM Cloud, you can add a local image of the VMDK format to a vCenter. Then, you can synchronize the vCenter image between the local client and the remote client by synchronizing data. Both system images and volume images can be added.

Event Message

Event Message allows you to check vCenter alarm messages, such as the message description, type, the vCenter from which the event message is sent, triggered user, target, and date.

- The UI can display up to 300 event messages. You can set a time range to check alarm messages within the time range via the time adjustment button at the upper left.
- You can choose to display alarm message count for each page via display count button at the lower right. Optional value: 10, 20, 50, and 100. You can turn pages by clicking the left/right arrow button.

2.2.1.9 Hybrid Cloud Management

Alibaba Cloud Hybrid Cloud Management provided by NexaVM Cloud integrates the simple, strong, scalable, and smart (4S) features of NexaVM Cloud Private Cloud and the advanced, secure, and

stable features of Alibaba Cloud Public Cloud. It is a hybrid cloud management solution that seamlessly integrates **cloud services and terminals**, interconnecting the control panel and data panel.

Concepts

- NexaVM Cloud Alibaba Cloud Hybrid Cloud Management provides the following cloud computing products of Alibaba Cloud:
  - ECS Instance: An elastic compute service (ECS) instance is a VM instance created on Alibaba Cloud.
  - Disk: A disk provides storage space for an ECS instance created on Alibaba Cloud.
  - Image: An image is a template file that is used to create ECS instances. Images are categorized into custom images and Alibaba Cloud images.
  - Security Group: A security group provides security control services for ECS instances on the L3 network. It filters the inbound or outbound packets of ECS instances based on security rules.
  - VPC: A virtual private cloud (VPC) is a private network dedicated for ECS instances created on Alibaba Cloud.
  - EIP: An elastic IP address (EIP) is an IP address in Alibaba Cloud public networks. You can attach EIPs to ECS instances so that the ECS instances can access public networks by using the EIPs.
- VPN: Establishes a site-to-site IPsec VPN channel to enable communications between private networks in a local data center and Alibaba Cloud VPC. This section includes:
  - VPN Gateway: A virtual private network (VPN) gateway establishes a secure connection between a local data center and Alibaba Cloud VPC by using an encrypted channel.
  - VPN Customer Gateway: A VPN customer gateway provides services for a local data center .
  - VPN Connection: A VPN connection is an encrypted communication channel established between a VPN gateway and VPN customer gateway.
- Express Connect: Express Connect uses physical circuits (electric cables or optical fibers leased from operators) to connect local data centers with Alibaba Cloud access points and Alibaba Cloud VPC. This way, private networks on Alibaba Cloud and in local data centers can communicate with each other in a fast, stable, and secure manner. This section includes:

- Router Interface: A router interface is a virtual device that is used to establish communication channels and control their status.
- Virtual Border Router: A virtual border router (VBR) is virtualized from a physical switch port on the access point of Alibaba Cloud. It forwards the data on the physical circuit to Alibaba Cloud VPC.

- Alibaba Cloud NAS: Alibaba Cloud NAS is a network-attached file storage service. It provides highly reliable and available distributed file systems that can be accessed by using standard file access protocols. In addition, Alibaba Cloud NAS is scalable in storage space and performance and can be managed in a namespace while shared with multiple users. NexaVM Cloud seamlessly integrates with Alibaba Cloud NAS. You can add primary storage of the **AliyunNAS** type on NexaVM Cloud Private Cloud so as to use the distributed storage independently deployed on Alibaba Cloud. This section includes:
  - File System: A file system is a backend storage system used for Alibaba Cloud NAS primary storage. Before you add an AliyunNAS primary storage, you need to add an NAS file system.
  - Permission Group: A permission group is an allowlist of IP addresses or IP ranges which can access file systems according to specified permission rules.
- Data Center: Data centers are resources corresponding to Alibaba Cloud regions and zones. These resources include:
  - Region: A region is a physical data center. A region in NexaVM Cloud Hybrid Cloud corresponds to a region in Alibaba Cloud.
  - Zone: A zone is a physical area in a region that is independent from other zones in the region in terms of electricity and network supplies.
- Setting: NexaVM Cloud Hybrid Cloud provides the following basic settings:
  - AccessKey Management: An AccessKey pair is an identity credential that has access to APIs of Alibaba Cloud or Private Alibaba Cloud. It has full access to the Cloud. An AccessKey pair consists of AccessKey ID and AccessKey secret.
  - Hybrid Cloud Settings: Hybrid cloud settings allow you to configure settings that take effect on the whole platform.

Physical Deployment

NexaVM Cloud Hybrid Cloud uses an in-process micro-service architecture and does not introduce a new module. NexaVM Cloud management nodes need to access the Internet so that they can call Alibaba Cloud Public Cloud APIs.

Physical connection-based deployment: uses physical connections to establish **local-remote** inter-connected networks, thereby connecting a local data center with Alibaba Cloud Public Cloud.

Figure 2-21: Physical Connection-based Deployment

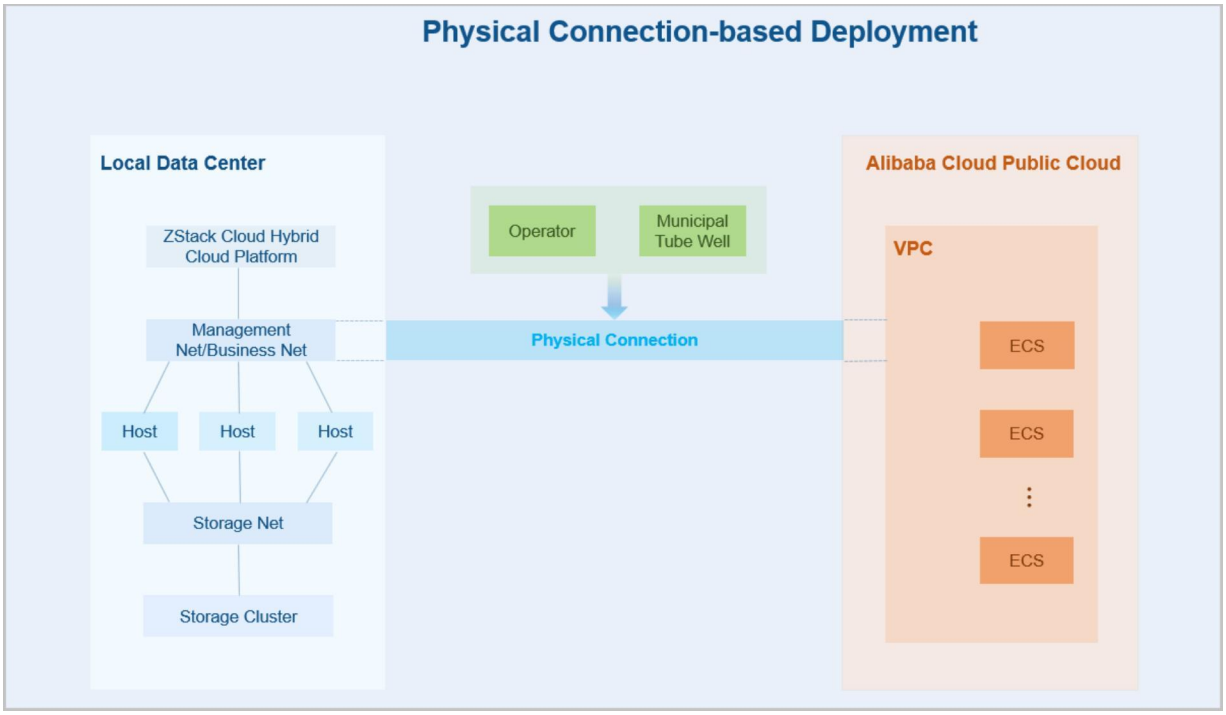

Architecture

NexaVM Cloud Hybrid Cloud includes the following sections:

- **Identity Authentication**:

  Alibaba Cloud AccessKey: integrates Resource Access Management of Alibaba Cloud Public Cloud / Private Cloud. A user authorized with an Alibaba Cloud AccessKey pair can access remote resources on Alibaba Cloud.

Figure 2-22: Identity Authentication

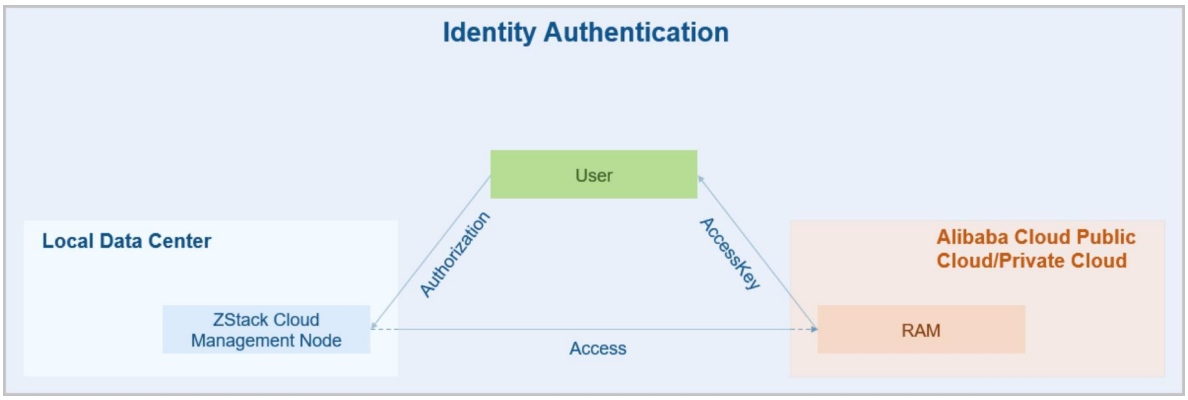

- **Network Interconnection**:

  You can use IPsec tunnels or Alibaba Cloud Express Connect to connect local Private Cloud with Alibaba Cloud Public Cloud. This way, **local-remote** L3 networks can access each other. The **Local-remote** network interconnection is the foundation of NexaVM Cloud Hybrid Cloud.

  NexaVM Cloud Hybrid Cloud allows you to use IPsec tunnels or Alibaba Cloud Express Connect to establish interconnected networks.

Figure 2-23: IPsec Tunnel

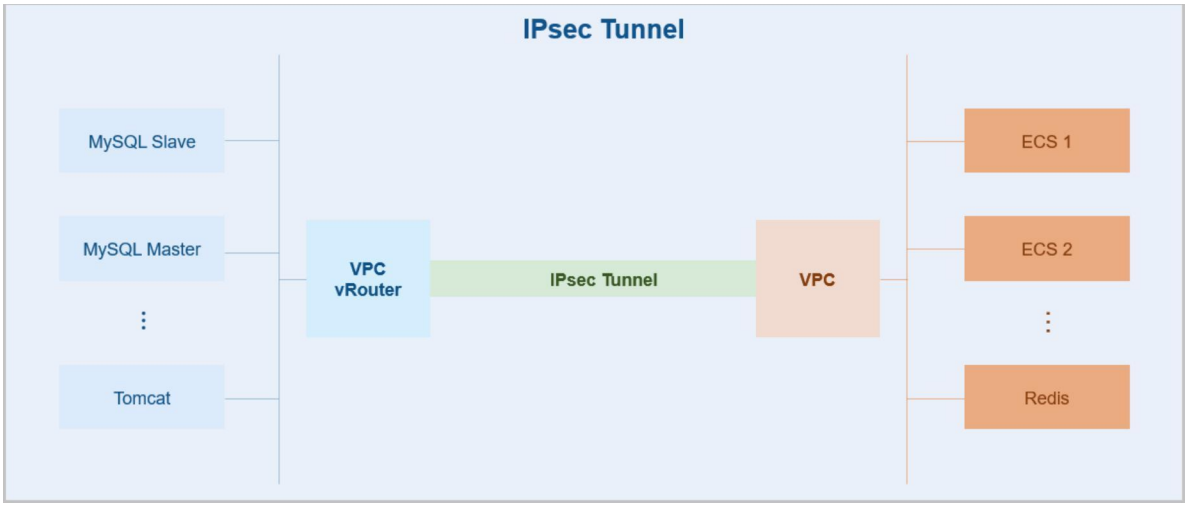

Figure 2-24: Alibaba Cloud Express Connect

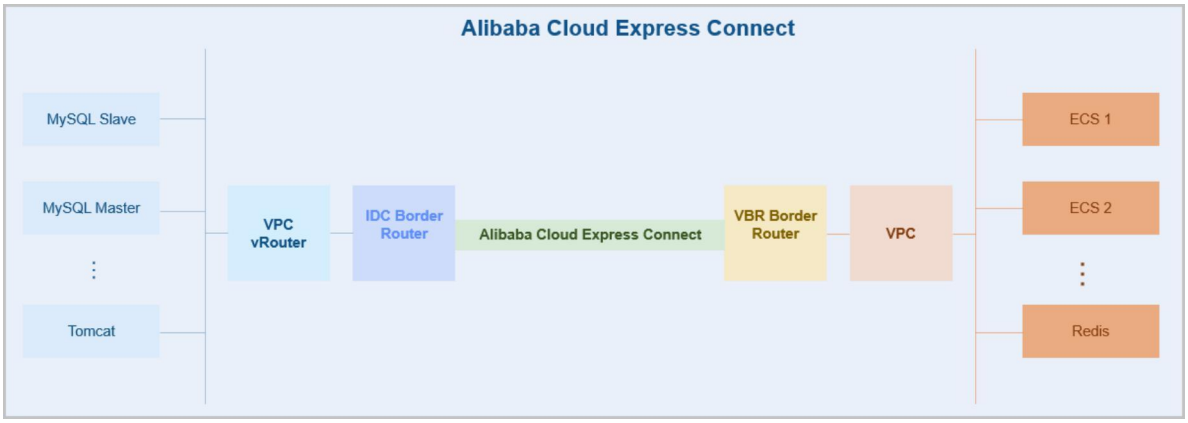

- **Resource Management**:

  You can authorize a RAM user to manage Alibaba Cloud Public Cloud resources, including ECS instances, VBR, VPC, and virtual switches.

  Resource Management

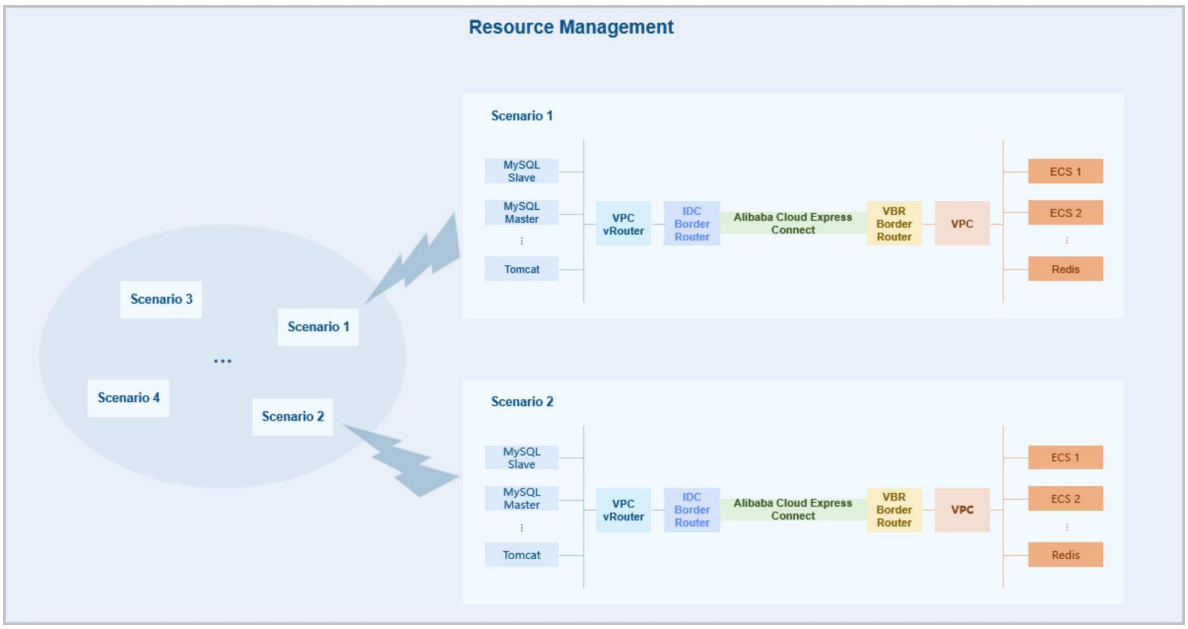

- **Business Implementation**:

  The identity authentication, network interconnection, and resource management mechanisms help establish a flexible and elastic business system architecture. After the hybrid cloud platform is established, you can deploy flexible and multi-dimensional business modes on it.

Characteristics

NexaVM Cloud Hybrid Cloud have the following characteristics:

- Seamless integration:

  NexaVM Cloud Hybrid Cloud seamlessly integrates Alibaba Cloud Public Cloud. Combined with the benefits of NexaVM Cloud Private Cloud, it provides users a platform to manage both public clouds and private clouds in a unified way.

- Seamless upgrading:

  NexaVM Cloud Hybrid Cloud allows seamless upgrading without affecting business continuity.

- Easy to use:

  NexaVM Cloud Hybrid Cloud seamlessly integrates **cloud services and terminals** in a unified cloud platform. You can easily manage local private clouds and access resources on the public cloud as needed.

Scenarios

- Data backup on the Cloud

  Financial, medical and some other industries have a high requirement for the compliance of long-term data storage. However, backing up data in local data centers is relatively risky, cost-consuming, and hard for O&M. To deal with these problems, NexaVM Cloud Hybrid Cloud helps you back up the data to the Cloud, providing you with a stable data storage service at a lower cost.

- Data storage on Cloud

  Enterprises and institutions need to store large amounts of data. In these scenarios, you can use NexaVM Cloud Hybrid Cloud to store data on Cloud. This solution lowers your investment and management costs and allows data access from multiple regions and zones.

- Data migration on Cloud

  High data negotiability is important to some enterprises and institutions whose works are finished based on multi-regional cooperation. In these scenarios, you can use NexaVM Cloud Hybrid Cloud to migrate data to Cloud, thus ensuring a stable data transmission and data integrity.

2.2.2 Platform O&M

2.2.2.1 Network Topology

A network topology visualizes the network architecture of the Cloud. It allows for efficient planning, management, and improvement of network architecture. Network topologies can be categorized into global topologies and custom topologies.

2.2.2.2 Cloud Monitoring

2.2.2.2.1 Management Node 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.

2.2.2.2.2 Performance Analysis

Performance Analysis displays the performance metrics of key resources monitored externally or internally in the Cloud. You can view the performance analysis or export the analysis report as needed to improve the O&M efficiency.

The key resources include: VM instance, host, VPC vRouter, backup storage, L3 network, and virtual IP address (VIP).

- **VM instance**: Displays the information such as the name, state, CPU utilization, memory utilization, disk read/write speed, NIC in/out speed, owner, and supported operation (stop VM instance) of a VM instance. You can customize the columns and choose to export the average, maximum, or minimum values of the metrics as needed.
- **VPC vRouter**: Displays the information such as the name, CPU utilization, memory utilization, disk read speed, disk write IOPS, NIC in rate, NIC in packets, NIC in errors rate, and owner of a VPC vRouter. You can customize the columns and choose to export the average, maximum, or minimum values of the metrics as needed.
- **Host**: Displays the information such as the CPU utilization, memory utilization, disk read/write speed, disk size, and NIC in/out speed of a host.
- **Backup storage**: Displays the information such as the name and available capacity percent of a backup storage.
- **L3 network**: Displays the information such as the name, used IPs, used IP percent, available IPs, and available IP percent of an L3 network.
- **VIP**: Displays the name, inbound/outbound traffic, inbound/outbound packet rate, and owner of a virtual IP address.

2.2.2.2.3 Capacity Management

Capacity Management visualizes the capacities and usages of key resources in the Cloud. You can use this feature to improve O&S efficiency.

Capacity Management displays the physical capacities and usages of key resources by using cards, and displays Top 10 resource capacities and usages, providing you a commanding view of your resource usages and greatly improving O&S efficiency.

2.2.2.2.4 Monitoring and Alarm

The Monitoring and Alarm feature monitors time-series data and events and sends alarm messages to specified endpoints by using SNS. Resource alarms, event alarms, and extended alarms are supported. The supported endpoints include the system, emails, DingTalk, HTTP applications, text messages, and Microsoft Teams. For some resource alarms, you need to install the agent before they can work as expected.

Concepts

- **Monitoring System**:

  A monitoring system provides the following features:

  - Monitor the following two types of time-series data:
    - Resource utilizations such as CPU utilization of VM instances and memory utilization of hosts
    - Resource capacities such as available number of IP addresses and total number of running VM instances
  - Event collection: collects events predefined on the Cloud, such as host disconnection and VM HA enabling.
  - Alarm: triggers alarms on time-series data or events.
  - Audit: records all operations and allows queries.
  - Customization: allows you to customize alarms and message templates and use predefined alarm templates and resource groups.
    - The following three types of alarms are supported:
      - Resource alarm: triggers alarms on time-series data. For example, you can configure an alarm for VM instances. If the CPU utilization of a VM instance exceeds 80% by five consecutive minutes, send an alarm message to an email address.

- Event alarm: triggers alarms on events, also called event subscription. For example, you can configure an alarm for host disconnection. If a host is disconnected, an alarm message is sent to DingTalk.
- Extended alarm: receives alarm messages from message sources. For example, if a Ceph Enterprise storage pool is downgraded, an alarm message is sent to the system of the Cloud.

— A message template specifies the text template of a resource alarm message or event alarm message sent to an SNS system.

- A message template and message recovery template are provided by the system. If you do not create a template, the system uses the predefined templates.
- You can create multiple message templates and can set only one template as the default template. Messages are formatted by using the default template.
- You can use `${}` in a template to quote variables configured in an alarm or event.
- You can configure email, DingTalk, text message, and Microsoft Teams as an endpoint in a message template. Messages sent by using email, DingTalk, text message, or Microsoft Teams are sent in the specified format.

— A message source is used to take over extended alarm messages. If you configure alarms for message sources, extended alarm messages can be sent to various endpoints. This enables centralized management of alarm messages and improves O&M efficiencies. You can configure a message source to take over alarm messages of Ceph Enterprise.

— An alarm template is a template of alarm rules. If you associate an alarm template with a resource group, an alarm is created to monitor the resources in the group.

— A resource group consists of resources grouped based on your business needs. If you associate an alarm template with a resource group, the alarm rules specified by the template take effect on all the resources in the group.

- **SNS**:

  SNS sends alarm messages to the specified endpoints. The supported types of endpoints include the system, emails, DingTalk, HTTP applications, text messages, and Microsoft Teams.

  Endpoints:

- The system provides the system-type endpoint. If you associate an alarm with this endpoint , alarm messages will be displayed below the Recent Message button in the top right corner of the UI.
- You can also create an endpoint of the email, DingTalk, HTTP application, text message, or Microsoft Teams type.

Characteristics

NexaVM Cloud monitoring and alarm has the following characteristics:

- Provides rich metric items to comprehensively monitor and alarm the core resources as well as events of the Cloud platform.
- Supports types of endpoints including system, emails, DingTalk, HTTP application, text messages, and Microsoft Teams for subscription topics. You can choose an appropriate endpoint to receive alarm messages according to the actual situation.
- One alarm can monitor multiple resources at the same time.
- Emails, DingTalk, text messages, and Microsoft Teams support customized alarm message templates. You can set alarm message templates on demand and quickly locate key informatio n from alarm messages.
- Supports for creating a template of alarm rules. If you associate an alarm template with a resource group, an alarm is created to monitor the resources in the group.

Scenarios

The function of monitoring and alarm monitors the core resources and events of the Cloud platform and sets up an alarm receiving mechanism. When core resources are abnormal, the monitoring and alarm will make real-time responses according to the alarm level to help O&M personnel quickly locate and solve the problem.

Global Setting

- Monitoring data is retained locally for 6 months by default, and you can customize the monitoring data retention period in the basic settings as follows:

  On the main menu of NexaVM Cloud, choose **Settings** > **Global Setting** > **Basic**. Then, the **Basic** tab is displayed. You can set **Monitoring Data Retention Period**. Enter an integer between 1 and 12. Default: 6. Unit: month.
- Monitoring data is retained locally in a size of 50GB by default, and you can customize the monitoring data retention size in the basic settings as follows:

On the main menu of NexaVM Cloud, choose **Settings** > **Global Setting** > **Basic**. Then, the **Basic** tab is displayed. You can set **Monitoring Data Retention Size** based on your needs. Default: 50 GB.

- NexaVM Cloud supports receiving extended alarm messages. On the main menu, choose **Settings** > **Global Setting** > **Advanced**. Then, the **Advanced** tab is displayed. You need to turn on the **Extended Alarm Notification** switch to use the extended alarm function.

2.2.2.2.5 SNS

SNS sends alarm messages to the specified endpoints. The supported types of endpoints include the system endpoint, email, DingTalk, HTTP applications, short message service, and Microsoft Teams.

- The Cloud provides a system endpoint by default. If an alarm binds a system endpoint, you are prompted for alarm notifications displayed near the Messages button at the upper right in the UI
- You can also create an email, DingTalk, HTTP application, short message service, or Microsoft Teams endpoint as needed.

Email Endpoint

- Messages that send to topics will be sent to a specified email address via an email server.
- You can either create an SNS text template in advance or use a system template to send emails in a unified format.
- You need to add an email server in advance under the current zone, and make sure that the email server works properly.

DingTalk Endpoint

- Messages that send to topics will be sent to a specified DingTalk robot address via DingTalk. If you appoint members, alarm notifications will be sent to corresponding DingTalk members via phone numbers.
- You can either create an SNS text template in advance or use a system template to send alarm messages in a unified format.
- If you set an SNS text template in DingTalk, follow the Markdown syntax. Currently, DingTalk only supports a subset of Markdown syntax.

HTTP Application Endpoint

- Messages that send to topics will be sent to a specified HTTP address via HTTP POST.

- If the specified HTTP application cannot be accessed without a user name and password, enter accurately the user name and the password.

Short Message Service Endpoint

- Messages that send to topics will be sent to a specified phone number via the short message service.
- You can create an SNS text template in advance and set it as the default template to send short messages according to the template that you set.

Microsoft Teams Endpoint

- Messages that send to topics will be sent to a specified Microsoft Teams group by using the webhook method.
- You can either create an SNS text template in advance or use a system template to send the Microsoft Teams messages in a unified format.

2.2.2.3 One-Click Inspection

One-Click Inspection: Comprehensively inspects the health status of key resources and services of the Cloud and scores their healthiness based on the inspection results. In addition, the one-click inspection service provides O&M suggestions and inspection reports. One-click inspection is applicable to centralized O&M scenarios.

Figure 2-25: One-Click Inspection

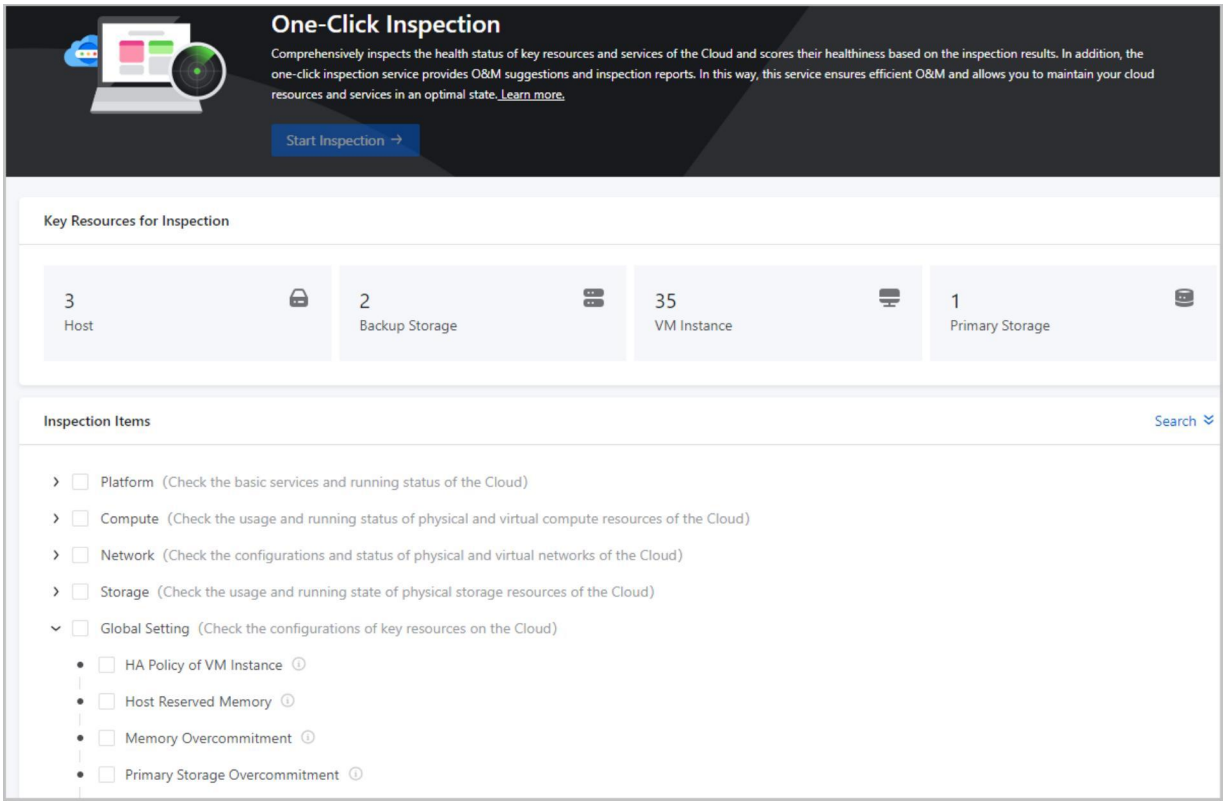

Concepts

- **Inspection Categories and Items**:

  One-Click Inspection provides five inspection categories, including platform, compute resources, network resources, storage resources, and global settings. You can use the service to inspect key resources and services of the Cloud, such as the management node, hosts, VM instances, backup storages, primary storages, physical and virtual NICs and networks, and licenses.

  - Platform: Check the basic services and running status of the Cloud.
  - Compute: Check the usage and running status of physical and virtual compute resources of the Cloud.
  - Network: Check the configurations and status of physical and virtual networks of the Cloud.
  - Storage: Check the usage and running status of physical storage resources of the Cloud.
  - Global Setting: Check the configurations of key resources of the Cloud.

  After you select items from certain inspection categories and launch inspection, related resources or services are inspected and their healthiness is scored. For more information about inspection items, see *Inspection Items*.

- **Inspection Results**:

  One-Click Inspection provides four inspection results, including Normal, Warning, Fault, Failed.

  - Normal: The inspected resources or services are in normal status. This result is marked with a green icon.
  - Warning: The health status of inspected resources or services is compromised, which may to some extent affect their performance and stability. This result is marked with a yellow icon .
  - Fault: The inspected resources or services are in critical condition and may seriously affect business operations. This result is marked with a red icon.
  - Failed: The inspection on related resources or services fails, which may seriously affect business operations. This result is marked with a grey icon.

- **Healthiness Scoring**:

  One-Click Inspection provides an in-built healthiness scoring mechanism for Cloud resources and services. It allows you to grasp the overall running status of the Cloud in a visualized way.

  Scoring on inspected resources/services: Scores resources and services based on the inspection results of related resource and service attributes.

  - If all attributes of a resource or service under inspection are in Normal status, the inspection result of the resource or service is Normal. The score is 100 points.
  - If one attribute of a resource or service under inspection is in Warning state and the other attributes are in Normal status, the inspection result of the resource or service is Warning. The score is 50 points.
  - If one attribute of a resource or service under inspection is in Fault or Failed state, the inspection result of the resource or service is Fault or Failed. The score is 0 points.

  Scoring on inspection items: Scores inspection items based on the inspection results of related resources and services.

  - If an inspection item does not belong to the Global Setting category, the inspection item is scored based on the following mechanism:
    - Score of Inspection Item = (Score of Resource 1 + Score of Resource 2 + …… + Score of Resource N)/(N*100)*100
    - For example, if an inspection item involves 3 resources, which are in Normal, Warning, and Fault/Failed status respectively, the scores of the three resources are 100, 50, and 0

points respectively. Then the score of the inspection item is (100 + 50 + 0)/(3*100)*100= 50 points.

- If an inspection item belongs to the Global Setting category, the inspection item is scored based on the following mechanism:
  - The score of the inspection item is the score of the involved global setting.
  - For example, if the inspection result of the involved global setting is Warning, the score of the global setting is 50 points. Then the score of the inspection item is 50 points.

Scoring on the Cloud: Scores the Cloud based on the scores of all inspection items.

- Score of the Cloud = (Score of Inspection Item 1 + Score of Inspection Item 2 + …… + Score of Inspection Item N)/(N*100)*100
- For example, if you select 3 inspection items, which is scored 100 points, 50 points, and 0 points respectively, then the score of the Cloud is (100 + 50 + 0)/(3*100)*100=50 points.

- **O&M Suggestions**:

  If resources and services are detected in Warning or Fault status, One-Click Inspection analyzes the hidden dangers and their effects on these resources and services, and provides suggestions on O&M. For more information, see *Inspection Items*.

- **Inspection Reports**:

  One-Click Inspection allows you to export PDF-formatted inspection reports. An inspection report summarizes platform configurations, resource status, and inspection results. It also provides details of all abnormal inspection items and corresponding O&M suggestions.

Benefits

One-Click Inspection has the following benefits:

- Comprehensive, customized, and efficient inspection capabilities: Provides five inspection categories that cover all key resources and services of the Cloud and allows you to select inspection items based on your business scenarios. After you launch an inspection, the inspection can be completed within a few minutes.
- Multi-layered scoring mechanism: The in-built three-layer mechanism of scoring on resources /services, inspection items, and the Cloud allows you to grasp the overall picture as well as details of the Cloud running status.
- Intelligent O&M suggestions: Provides risk analysis of resources and corresponding countermeasures, facilitating efficient O&M.

2.2.2.4 Message Log

2.2.2.4.1 Alarm Message

An alarm message is a message sent the time when an alarm is triggered.

2.2.2.4.2 Operation Log

An operation log is a chronological record of operations on the specified objects and their operation results.

2.2.2.4.3 Current Task

A current task is an ongoing operation performed in the Cloud. You can perform centralized management over ongoing operations.

2.2.2.4.4 Audit

Audit monitors and records all activities on the Cloud. You can use this feature to implement operation tracking, cybersecurity classified protection compliance, security analysis, troublesho oting, and automatic O&M.

2.2.2.4.5 Audit

Allows you to collect with one click the log data from the Cloud and various nodes on the Cloud generated in the specified time period and download the log data.

2.2.2.5 Backup Management

2.2.2.5.1 Backup Service

Backup management integrates multiple disaster recovery technologies such as incremental backup and full backup that are suitable for multiple business scenarios. You can implement local backup and remote backup based on your business needs.

Backup Service is a separate feature module. To use this service, purchase both the Base License and the Plus License of Backup Service. The Plus License cannot be used independently.

Typical Backup Scenarios

Backup Service can be applied to the following three typical scenarios: local backup, remote backup, and Public Cloud backup.

Note:

If you have the Tenant Management Plus license at the same time, the project members (project managers, project admins, and general project members) can perform local backup for VM instances and volumes in the project.

- **Local Backup**

  A local ImageStore backup storage can act as the **Local Backup Storage** to store scheduled backup data of the local VM instances, volumes, and management node databases. Meanwhile, the seamless switchover between the primary local backup storage and the secondary local backup storage is supported, which effectively ensures your business continuity.

  If your local data is mistakenly deleted, or data in the local primary storage is damaged, you can recover the backup data from the local backup storage, as shown in *Local Backup Scenario-1*

  Figure 2-26: Local Backup Scenario-1

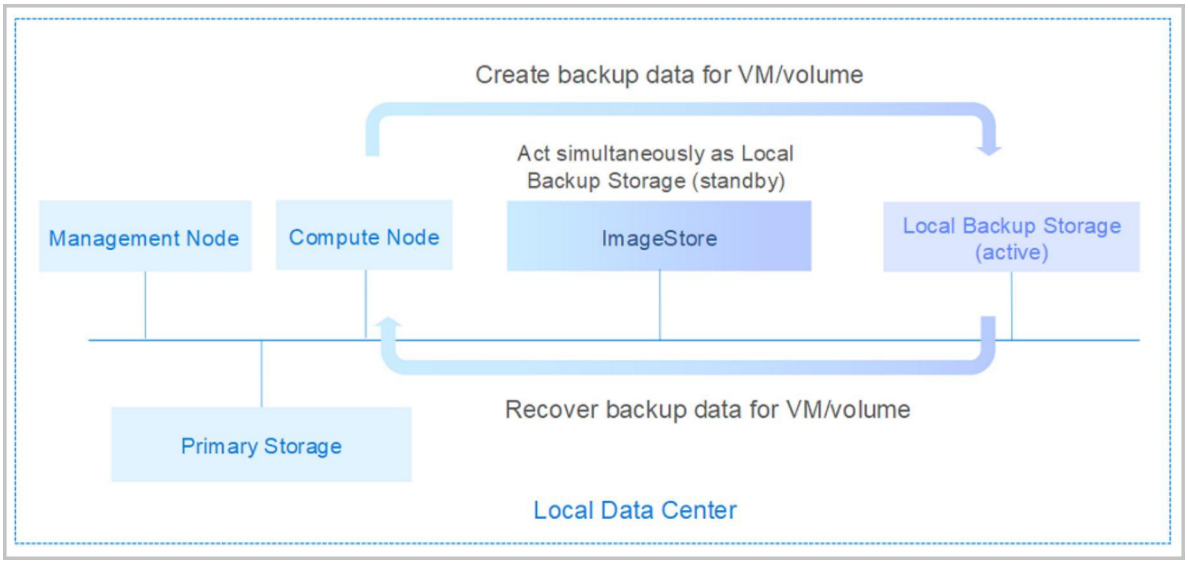

  If you encounter a disaster in your local data center, you can rely totally on your local backup storage to rebuild your data center and recover your business, as shown in *Local Backup Scenario-2*.

Figure 2-27: Local Backup Scenario-2

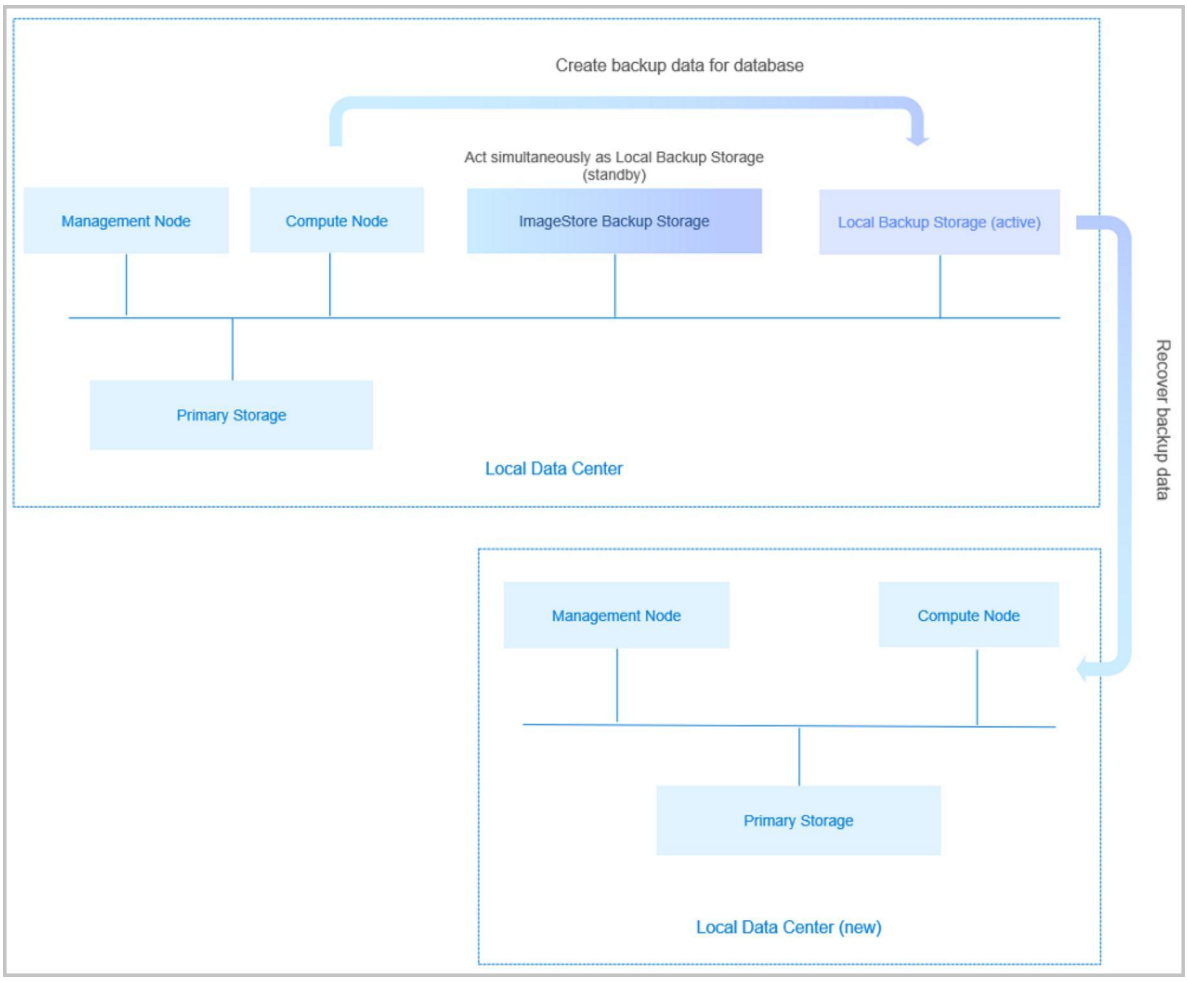

- **Remote Backup**

A storage server in a remote data center can act as the **Remote Backup Storage** to store the scheduled backup data of the local VM instances, volumes, and databases. The backup data needs to be synchronized to the remote backup storage from the local backup storage.

If your local data is mistakenly deleted, or data in the local primary storage is damaged, you can recover the backup data from the remote backup storage, as shown in *Remote Backup Scenario-1*

Figure 2-28: Remote Backup Scenario-1

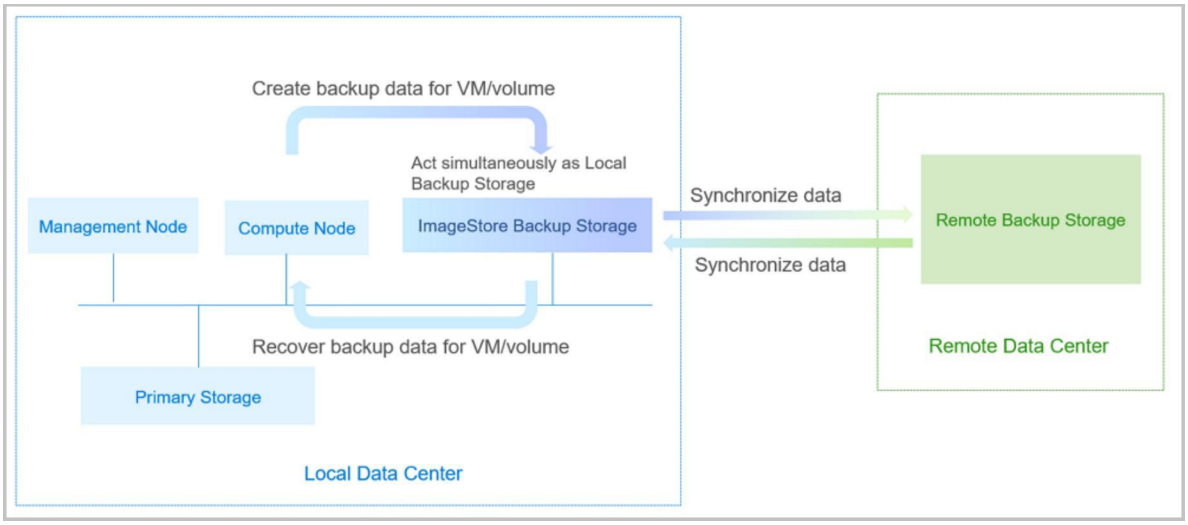

If you encounter a disaster in your data center, you can rely totally on your remote backup storage to rebuild your data center and recover your business, as shown in *Remote Backup Scenario-2*.

Figure 2-29: Remote Backup Scenario-2

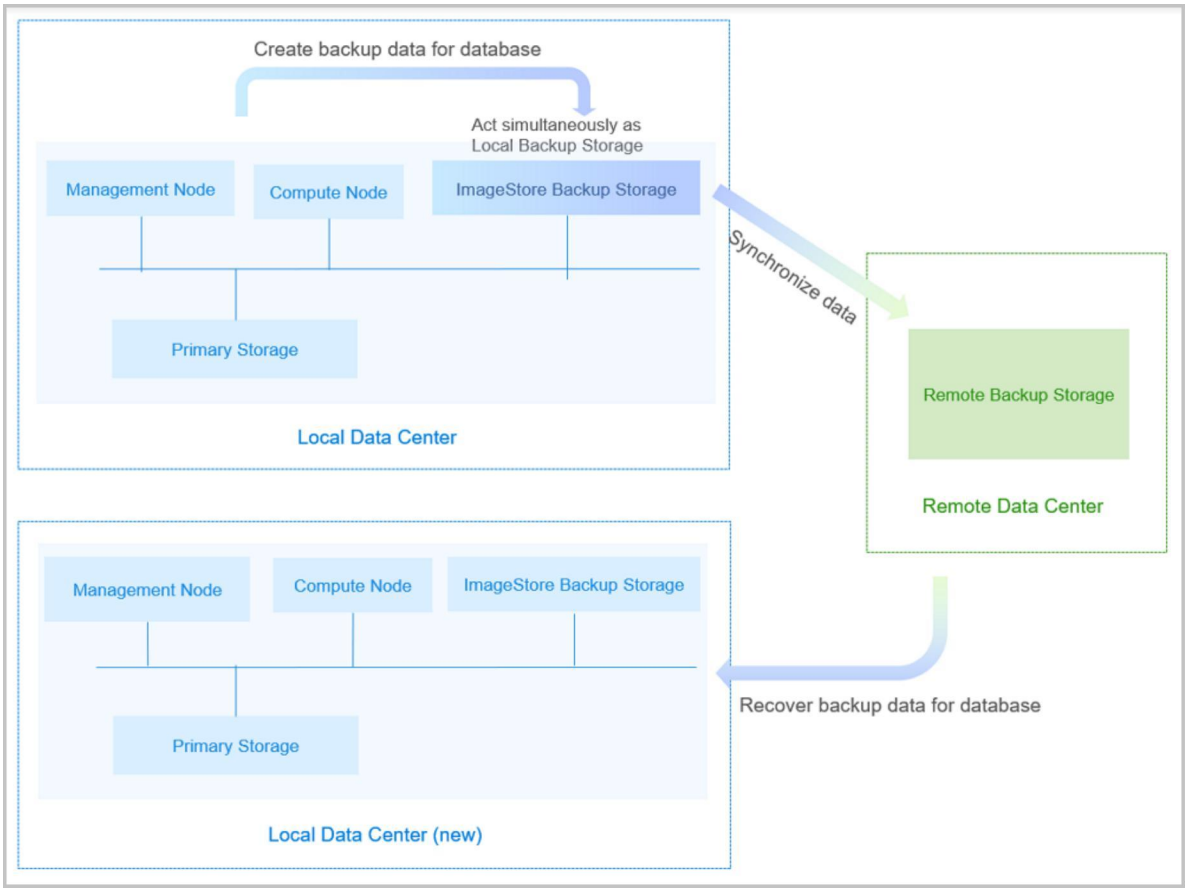

- **Public Cloud Backup**

The storage server in the Public Cloud can act as the **Public Cloud Backup Storage** to store the scheduled backup data of the local VM instances, volumes, and databases. The backup data can be synchronized to the Public Cloud backup storage from the local backup storage.

If your local data is mistakenly deleted, or data in the local primary storage is damaged, you can recover the backup data from the Public Cloud backup storage, as shown in *Public Cloud Backup Scenario-1*.

Figure 2-30: Public Cloud Backup Scenario-1

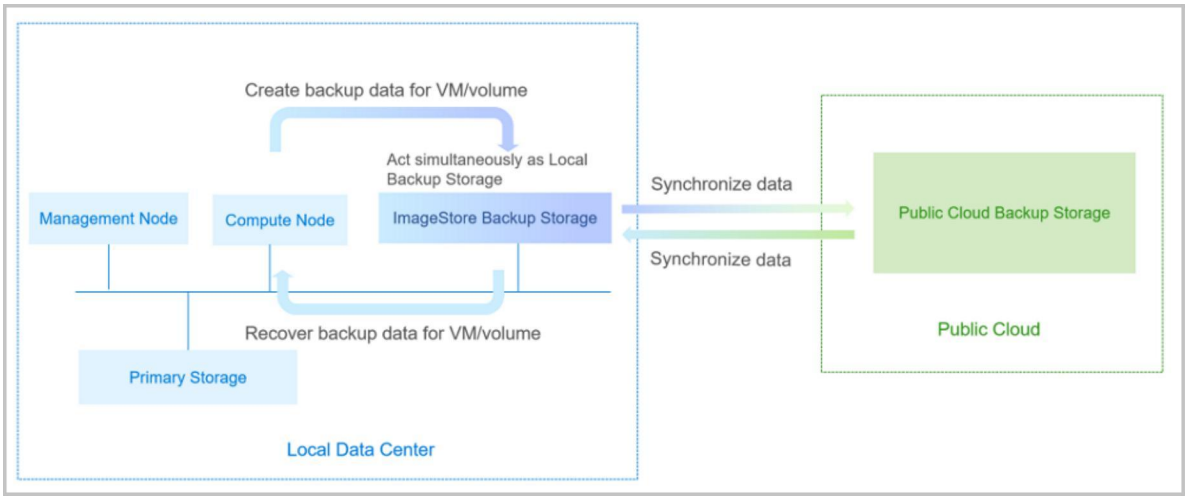

If you encounter a disaster in your data center, you can rely totally on your Public Cloud backup storage to rebuild your data center and recover your business, as shown in *Public Cloud Backup Scenario-2*.

Figure 2-31: Public Cloud Backup Scenario-2

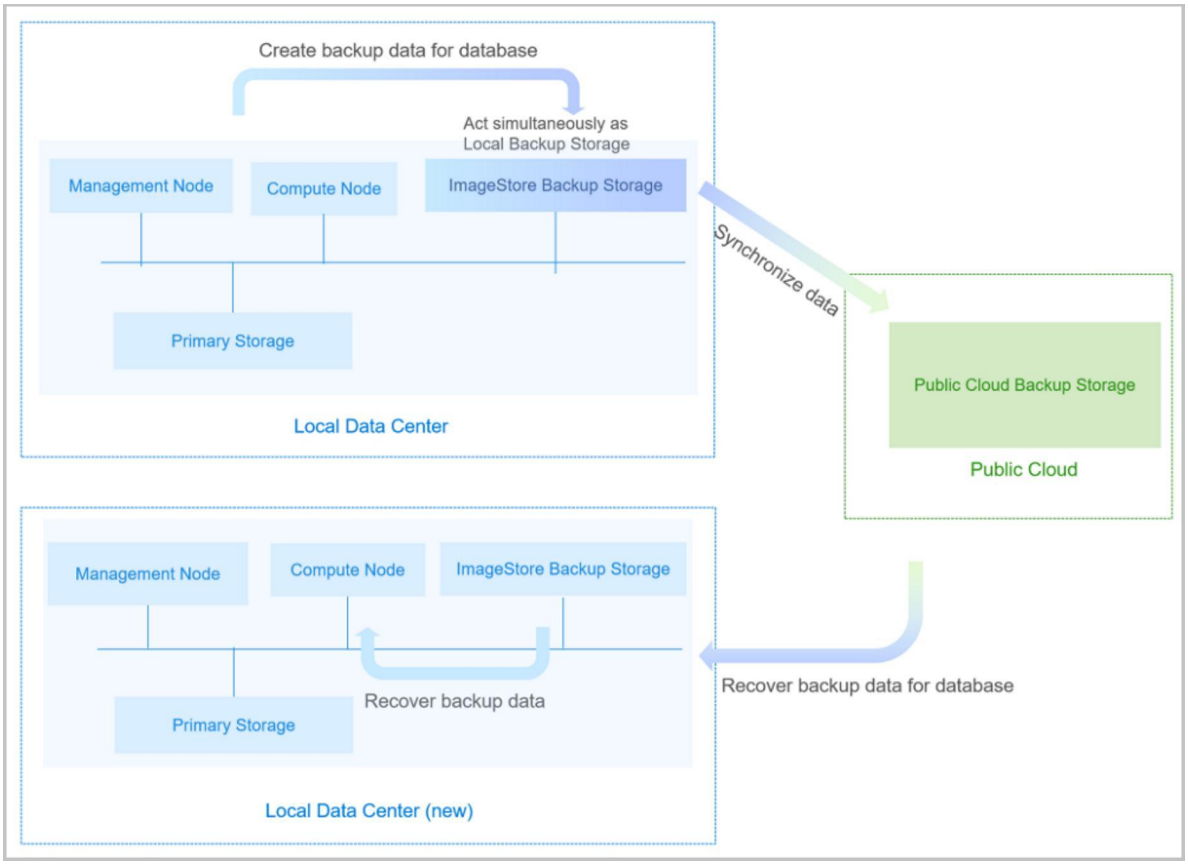

2.2.2.5.1.1 Backup Job

You can create backup jobs to back up VM instances, volumes, and databases on your local data center to specified local storage servers and sync the backup data to specified remote backup servers or backup servers on the public cloud.

2.2.2.5.1.2 Local Backup Data

The local backup data is the data of local VM instances, volumes, and databases backed up in the local backup server. You can manage the local backup data on the **Local Backup Data** page.

- You can restore the backup data to your local server or synchronize the data to a remote backup server.
- When you restore a database, refreshing the browser will make the UI display improperly without affecting the restoring process.
- With the Backup Service, you can back up or restore the data stored in the management node database. Note that the operation logs and monitoring information cannot be backed up or restored.

2.2.2.5.1.3 Local Backup Server

A local backup server is located at the local data center and is used to store local backup data of VM instances, volumes, and databases.

- You can use the ImageStore deployed in the local data center as the local backup server.
- You can also deploy a new local backup server.
- You can add more than one local backup server.
- If you specify more than backup server for a backup job, these backup servers can work in the active-standby mode.
- You can clean up invalid and expired data that is completely deleted to free up storage space.
- You can view the backup data backed up to the local backup server on the details page.

2.2.2.5.1.4 Remote Backup Server

A remote backup server is located at a remote data center or a Public Cloud and is used to store remote backup data of VM instances, volumes, and databases.

- The backup data can only be synchronized from a local backup server to the remote backup server.
- You can add only one remote backup server to the Cloud.
- You can view the backup data backed up to the remote backup server on the details page.
- Before you can restore the remote backup data of VM instances and volumes to your local server, synchronize the data from the remote backup server to a local backup server first.
- The remote backup data of databases can be restored the your local server directly.

2.2.2.5.2 Continuous Data Protection (CDP)

Continuous Data Protection (CDP) provides second-level and fine-grained continuous backups for important business systems in VM instances, allowing users to restore VM data to a specific time state, and retrieve files without restoring the system. NexaVM Cloud provides automated solutions on resuming all applications if a hardware or operating system failure occurs.

The Continuous Data Protection (CDP) service is a separate feature module. To use this service, purchase both the Base License and the Plus License of Continuous Data Protection (CDP). The Plus License cannot be used independently.

Concepts

- CDP task: You can create a CDP task to continuously back up your VM data to a specified backup server to achieve continuous data protection and recovery.
- Recovery task: A recovery task helps you quickly restore data by specifying a CDP task and recovery point, and allows you to view the recovery progress and logs in a more friendly way.
- CDP data: The backup data generated from continuous data protection on VM instances is stored in local backup servers.
- Recovery point: A recovery point is a data point generated during continuous data protection. A recovery point corresponds to a data record within the recovery point interval specified by the user.
- Locked recovery point: You can lock or unlock a recovery point as needed. After a recovery point is locked, data of the recovery point will not be automatically cleared or deleted.

How CDP Works

NexaVM Cloud provides block-level continuous data protection. You can restore CDP data according to a specified time point.

- **CDP backup**

  After continuous data protection is performed on a VM instance, the CDP server first performs a full copy of the VM data, continuously captures I/O data changes, and timestamps, and saves the I/O data of each change, thereby achieving continuous data backup.

Figure 2-32: CDP Backup

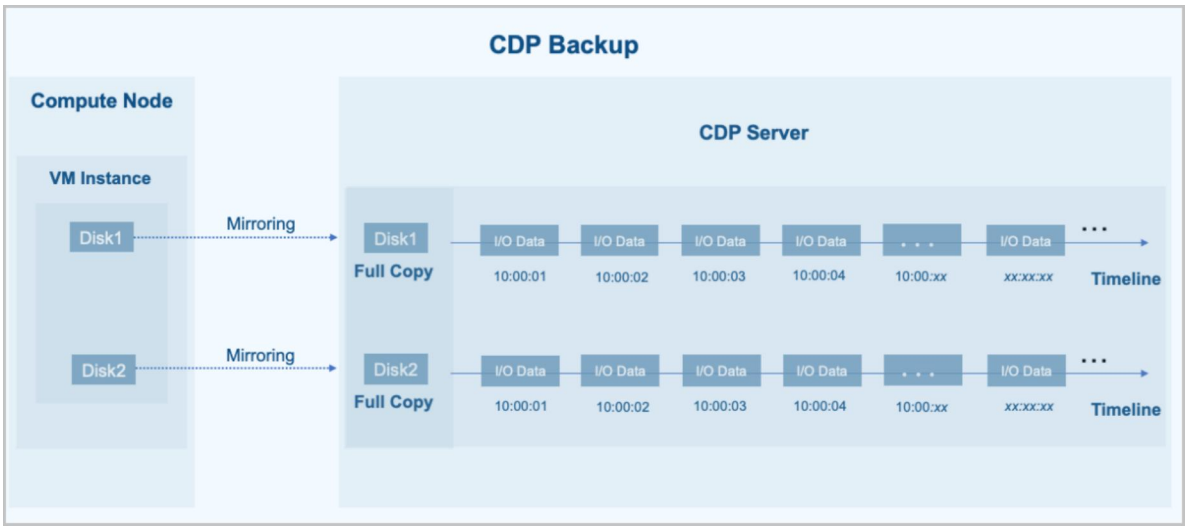

- **CDP recovery**

When data recovery is performed, the CDP server exposes the CDP data as a block device and then restores the I/O data at a specified time to the disk or file system of a primary storage.

Figure 2-33: CDP recovery

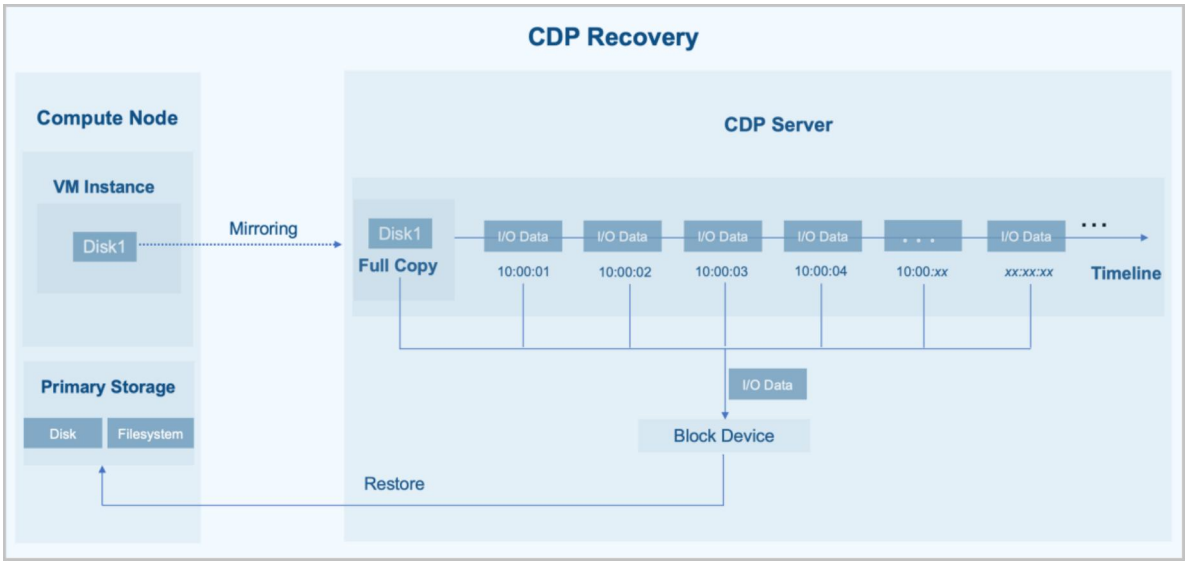

Advantages

- Simple:
  - The software-defined solution we provide is hardware-independent and scalable.
  - You can preview the backup files without restoring the system. Supported file formats: pictures (.png, .jpg, .bmp, .gif, etc.), PDF files, and text files (equal to or smaller than 10 MB).
  - When you create a CDP task for the first time, the Cloud intelligently recommends the desired capacity required by a CDP task based on an algorithm, helping you to plan the backup space reasonably.
  - You can restore data through a wizard-style process.
- Strong:
  - Agentless backup: To use the CDP service, you do not need to install agents for your VM instances or couple with other applications. This helps reduce the configuration complexity, lower the VM performance loss, and guarantee the business security.
  - Second-level RPO: Provides second-level fine-grained continuous data protection for VM instances.

- Instant recovery: Supports instant recovery with the RTO in seconds, which helps to ensure the business continuity.
- Primary storage support: The CDP service applies to VM instances in different primary storage scenarios, including local, NFS, SharedBlock, and Ceph primary storages.

- Flexible:
  - Flexible RPO settings: Provides second/minute-level RPO settings.
  - Multiple recovery levels: Supports entire recovery and file-level recovery.
    - Entire recovery: You can restore data to the original VM instance or to a newly-created VM instance.
    - File-level recovery: You can retrieve files without restoring the system. Both Windows and Linux file system formats are supported.
  - Flexible data display and search:
    - The CDP data page displays hourly data changes of a VM instance, which provides a reference for backup capacity planning.
    - The CDP data page also provides a recovery point calendar, which identifies the dates with recovery points with colors. This helps you locate a recovery point quickly.

- Reliable:
  - Unified O&M: You can view the critical CDP information on the CDP overview page, including the CDP task status, recovery task status, backup server usage, and unread CDP alarms.
  - You can restore CDP data to the original VM instance by creating a volume for the VM instance. All of the volumes before recovery can be retained and attached to the VM instance again, which ensures the data security to the maximum extent and facilitates post-fault analysis.
  - The "Create VM instance" recovery policy allows you to create a new VM instance from the selected recovery point without affecting the original VM instance. You can finish the recovery after you confirm that the data is correct. This helps to meet the recovery drill requirements.
  - You can mark and lock recovery points to retain the recovery point data in long term.
  - The Cloud provides a list of recovery tasks, allowing you to view the recovery records and progress in a more friendly way.

- The RPO latency policy and alarm help to effectively relieve data transmission pressure of backup servers in heavy I/O scenarios.

Scenarios

The CDP service can be applied in the following scenarios:

- Continuous protection of critical business data

  You can use the CDP service to continuously back up and protect critical business data, such as banking system data and financial transaction data, minimizing data loss and ensuring your business continuity.

- Anti-virus and data recovery

  In case of virus invasion, such as a ransomware attack, you can use CDP to restore data to any point in time to reduce the loss caused by the virus.

Limits

Currently, if a VM instance has volumes not stored on the Ceph primary storage, you could not clone the VM instance or create snapshots or images for the VM instance during CDP.

2.2.2.5.2.1 CDP Dashboard

The overview page displays the critical CDP information on different cards, including the number and status of CDP tasks and recovery tasks, the CPU and memory utilization of backup servers , top 5 backup server usage, the total disk I/O of backup servers, and unread alarm statistics in recent 7 days.

On the main menu of NexaVM Cloud, choose **Platform O&M** > **Backup Management** > **CDP Service** > **CDP Dashboard**. Then, the **CDP Dashboard** page is displayed.

Figure 2-34: CDP Dashboard

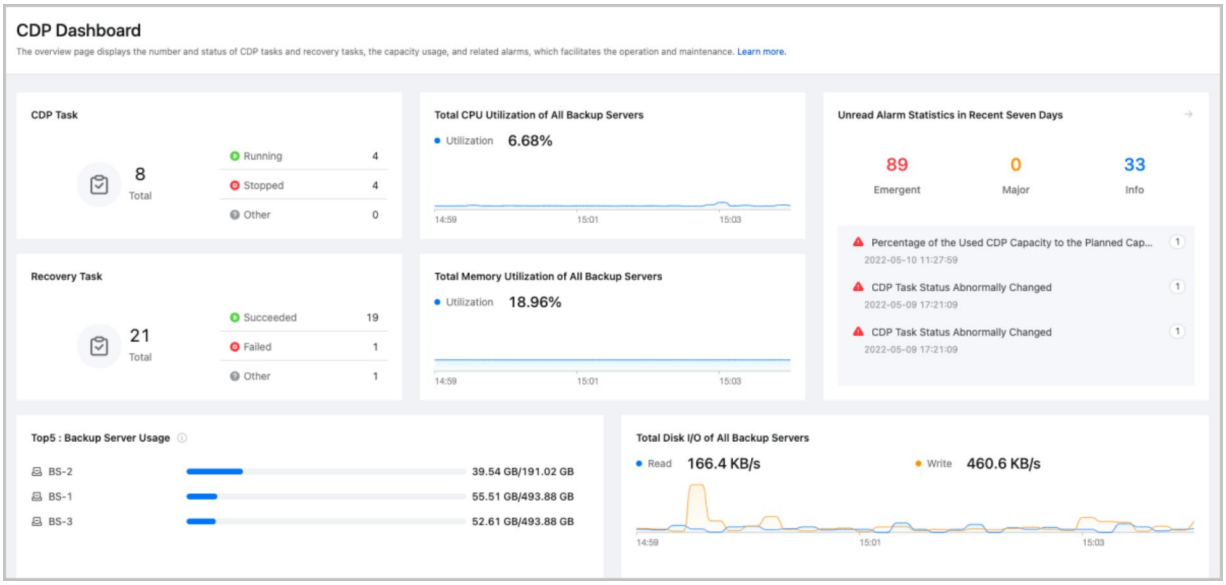

The cards are described as follows:

- CDP Task:
  - This card displays the number and status of CDP tasks in the Cloud.
  - Task status includes running, stopped, and other (starting, running, unknown, and failed).
  - You can click the number on the card to enter the CDP task page to view more information.
- Recovery Task:
  - This card displays the number and status of recovery tasks in the Cloud.
  - Task status includes succeeded, failed, and other (waiting, paused, recovering, canceling, and canceled).
  - You can click the number on the card to enter the recovery task page to view more information.
- Total CPU Utilization of All Backup Servers: This card displays the CPU utilization of all backup servers in the Cloud.
- Total Memory Utilization of All Backup Servers: This card displays the memory utilization of all backup servers in the Cloud.
- Top 5 Backup Server Usage:
  - This card displays the used capacity and total size of each backup server.
  - The usage of each backup server is displayed in descending order.

  - You can click the backup server name on the card to enter the details page of the backup server.
- Total Disk I/O of All Backup Servers: This page displays the disk I/O of all backup servers in the Cloud.
- Unread Alarm Statistics in Recent Seven Days:
  - This card displays unread alarm statistics in recent 7 days, including the emergency level, number of alarms, and alarm name.
  - You can click the **More** icon in the upper right corner to enter the alarm message page.
  - You can view and handle the alarm messages and copy the alarm details.
  - Alarm messages that you already read are not displayed here again.

2.2.2.5.2.2 CDP Task

You can create a CDP task to continuously back up your VM data to a specified backup server to achieve continuous data protection and recovery.

- Before you can use the CDP service, add a local backup server to the Cloud first.
- You can create CDP tasks to continuously back up your VM data to a specified backup server to achieve continuous data protection.
- You can create CDP tasks in bulk for multiple VM instances. The Cloud support only one VM instance per CDP task.
- You can perform entire VM backup without installing an agent for your VM instances.
- The Cloud performs a full backup on the VM instances immediately after you create CDP tasks.
- The Cloud provides second-level fine-grained continuous data protection for VM instances.
- The Cloud recommends the desired capacity required by a CDP task based on an algorithm when you create a CDP task for the first time, helping you to plan the backup space reasonably .
- The CDP service applies to VM instances in different primary storage scenarios, including local , NFS, SharedBlock, and Ceph primary storages.
- You can manage the lifecycle of CDP tasks, such as creating, enabling, disabling, and deleting CDP tasks.
- You can modify the protection policy of a CDP task, including the recovery point interval, regular backup frequency, recovery point retention policy, and the backup rate when the CDP task is disabled.
- You can modify task running policy to adjust the desired size and RPO policy for the CDP task.

- You can view the creation progress of a CDP task.
- The Cloud provides CDP task resource alarms and event alarms and allows you to create these alarms.

2.2.2.5.2.3 CDP Data

The backup data generated from continuous data protection on VM instances is stored in local backup servers. You can manage CDP on the **CDP Data** page

- You can back up CDP data on a local backup server.
- The Cloud displays the CDP status in charts and tables and allows you to view the details by specifying a time span.
- The Cloud displays hourly data changes so that you plan the backup capacity more reasonably .
- The Cloud provides a recovery point calendar, which identifies the dates with recovery points with colors and helps you to locate recovery points quickly.
- You can lock recovery points. After a recovery point is locked, data of the recovery point will not be automatically cleared or deleted.
- The Cloud provides recovery point list and locked recovery point list and allows you to view the details by specifying a time span.
- The Cloud supports fast recovery based on selected recovery points (including locked recovery points).
- The Cloud supports instant recovery with a minimum RTO in seconds.
- The Cloud supports entire restoration and file-level restoration.
  - Entire restoration allows you to restore data to the original VM instance or to a newly-created VM instance.
    - Restore data to a newly-created VM instance:
      - Allows you to create a VM instance from the selected recovery point without affecting the original VM instance.
      - The newly created VM instance will quickly start up for business recovery.
    - Restore to the original VM instance:
      - Allows you create new volumes or overwrite current volumes.
        - Create new volumes: This method allows you to retain and attach volumes before recovery to the VM instance to ensure data security.

        - Overwrite current volumes: This method will overwrite the original data in the VM instance and keep the snapshots in the current volumes.
    - During data restoration, the original VM instance will quickly start up for business recovery.
- File-level restoration allows you to retrieve files without restoring the system. Both Windows and Linux file system formats are supported. Supported file format include picture, text, and PDF.

- Allows you to clear CDP data, which will delete all the CDP data of the VM instance, including the locked recovery points. The Cloud performs full backup for the VM instance the next time the CDP task is enabled.

2.2.2.5.2.4 Recovery Task

A recovery task helps you quickly restore data by specifying a CDP task and recovery point, and allows you to view the recovery progress and logs in a more friendly way.

- The Cloud provides a list of recovery tasks, allowing you to view the recovery records and progress in a more friendly way.
- The CDP service applies to VM instances in different primary storage scenarios, including local , NFS, SharedBlock, and Ceph primary storages.
- The Cloud supports instant recovery with a minimum RTO in seconds.
- The Cloud allows you to restore data to the original VM instance or to a newly-created VM instance.
    - Restore data to a newly-created VM instance:
        - Allows you to create a VM instance from the selected recovery point without affecting the original VM instance.
        - The newly created VM instance will quickly start up for business recovery.
    - Restore to the original VM instance:
        - Allows you create new volumes or overwrite current volumes.
            - Create new volumes: This method allows you to retain and attach volumes before recovery to the VM instance to ensure data security.
            - Overwrite current volumes: This method will overwrite the original data in the VM instance and keep the snapshots in the current volumes.

- During data restoration, the original VM instance will quickly start up for business recovery.

- You can manage the lifecycle of recovery tasks, such as creating, enabling, disabling, and deleting recovery tasks.
- You can rerun a failed or canceled recovery task.
- You can cancel a task only during the recovery progress. After a task is canceled, intermediate data generated during the recovery process will not be retained.

2.2.2.5.2.5 Local Backup Server

A local backup server is located at the local data center and is used to store local CDP data.

- You can use the ImageStore deployed in the local data center as the local backup server.
- You can also deploy a new local backup server.
- You can add more than one local backup server.
- You can view the CDP data backed up to the local backup server on the details page.

2.2.2.6 Scheduled O&M

2.2.2.6.1 Scheduled Job

NexaVM Cloud provides two types of scheduled O&M resources: scheduled jobs and schedulers. These two types of resources are independent from each other. You can create schedulers and

scheduled jobs based on different rules, and associate or disassociate scheduled jobs with or from schedulers.

2.2.2.6.2 Scheduler

NexaVM Cloud provides two types of scheduled O&M resources: scheduled jobs and schedulers. These two types of resources are independent from each other. You can create schedulers and

scheduled jobs based on different rules, and associate or disassociate scheduled jobs with or from schedulers.

- A scheduled job defines that a specific action be implemented at a specified time based on a scheduler.

  - You can associate any available scheduled job with a scheduler.
  - You can select Disable, Enable, Attach, and Detach actions for a scheduled job based on your actual production environments.

  - If you delete a scheduler, the scheduled jobs associated with the scheduler will be disassociated. You can associate the scheduled jobs with other schedulers.
  - Operations triggered by scheduled jobs are all recorded by the Audit feature.
- A scheduler is used to schedule jobs. It is suitable for business scenarios that last for a long time.
  - A scheduler defines the implementation rules for a scheduled job.
  - A scheduler can be used for long-term operations, for example, creating snapshots at a specified interval for a VM instance.
  - If you delete a scheduler, the scheduled jobs associated with the scheduler will be disassociated. You can associate the scheduled jobs with other schedulers.
  - Operations triggered by schedulers are all recorded by the Audit feature.

2.2.2.7 Tag Management

A tag is used to mark resources. You can use a tag to search for and aggregate resources. Specifically, you can quickly locate the required resources by tag type and tag name.

- You can create tags with different colors, simple style, and brief description. You can also attach tags to resources and search resources by using tags. This will improve your search efficiency.
- You can search for the resources without tags by clicking the option "None" when you use tags to filter resources. This is convenient for maintenance operations.
- Two types of tag are available: admin tags and tenant tags.
  - Admin tags are created and owned by the administrator, and can be attached to VM instances, volumes, hosts, baremetal instances, and elastic baremetal instances.
  - Tenant tags are created and owned by tenants, and can be attached to VM instances and volumes.
- Currently, you can attach tags to or detach tags from VM instances, volumes, hosts, baremetal instances, and elastic baremetal instances.

Considerations

- Admin tags are created and owned by the administrator while tenant tags are created and owned by tenants.
- Tags created by tenants can only be attached to resources of the corresponding tenants, while admin tags can be attached to all of the resources on the Cloud.

- The administrator can detach or delete tenant tags.
- Tags in a project are owned by the project. Therefore, everyone in the project, including the project admin, project manager, and project member, can perform operations on these tags.
- Currently, tag owners cannot be changed.
- When you change a resource owner, all tenant tags attached to the resource will be detached. However, the admin tags are not affected.
- After the Cloud is upgraded seamlessly, the existing tags will be updated accordingly and displayed in the latest way. If an exception occurs, refresh your browser or create a new tag.

2.2.2.8 Migration Service

NexaVM Cloud provides the Migration Service, allowing you to migrate VM systems and data from other virtualization platforms to the current cloud. Currently, with the Migration Service, you can:

- Migrate VM instances from the vCenter that you took over to the current cloud. The supported vCenter versions include 5.5, 6.0, 6.5, 6.7, and 7.0. Note that the version of the vCenter server must be consistent with that of the ESXi host.
- Migrate VM instances from a KVM cloud platform to the current cloud.

Note:
If you took over vCenter 7.0, to ensure that the VM console can open properly, we recommend that you download the trusted root CA certificate when you log into vCenter.

Figure 2-35: V2V Migration

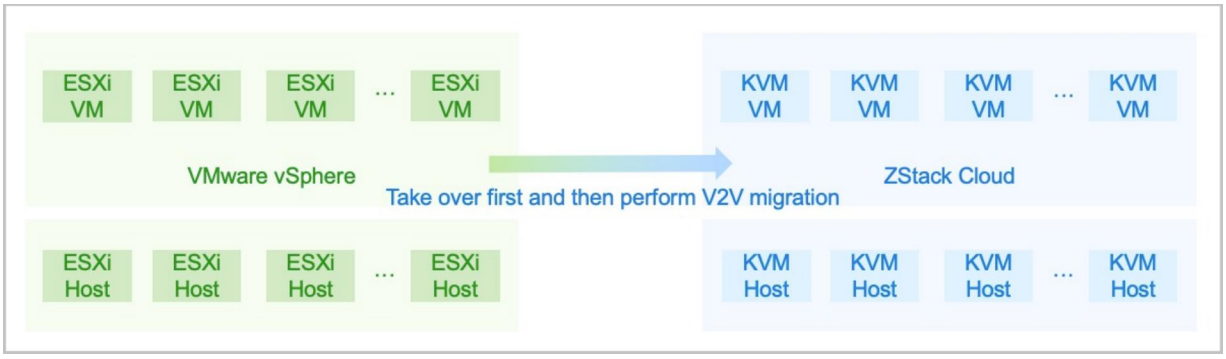

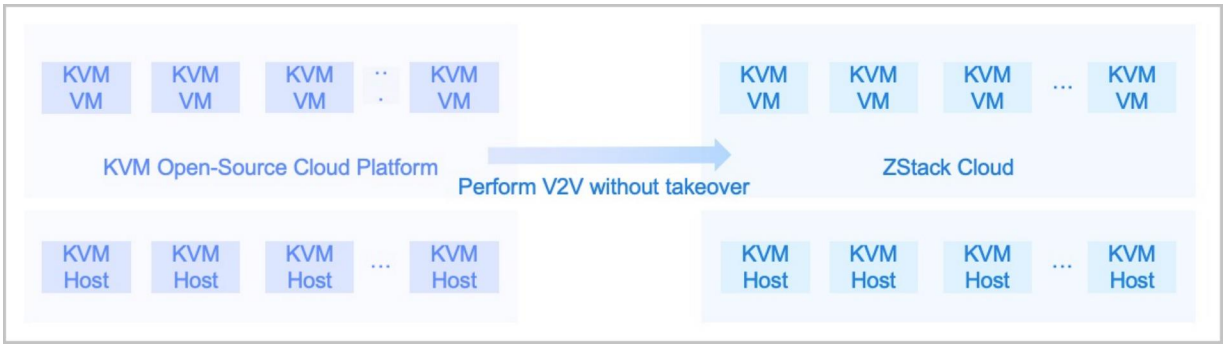

The Migration Service is a separate feature module. To use this feature, you need to purchase both the Base License and the Plus License of the Migration Service. The Plus License cannot be used independently.

Advantages of the Migration Service are as follows:

- Allows you to perform one-click V2V migrations for VM instances in bulk.
- Allows you to add a conversion host and create a V2V job and lets the Cloud do the rest.
- Allows you to configure an independent migration network and a network QoS for a conversion host to control transmission bottlenecks and improve migration efficiencies.
- Allows you to customize configurations for destination VM instances when you create a V2V job.
- Monitors and manages the entire migration process in the visualized, well-designed UI.

2.2.2.8.1 V2V Migration

Currently, you can migrate VM instances from a VMware cloud platform or a KVM cloud platform to the current cloud.

Source Cloud Platform: VMware

You can migrate VM instances from the vCenter you take over to the current Cloud by creating a migration task.

- Before migrations, perform **data synchronization** to manually synchronize the latest status of resources in the vCenter that you took over.
- You can perform bulk V2V migrations for VM instances, and customize configurations of the migrated VM instances.
- The supported vCenter versions include 5.0, 5.1, 5.5, 6.0, 6.5, 6.7, and 7.0. Note that the version of the vCenter server must be consistent with that of the ESXi host.

- The supported VM systems of the source vCenter include RHEL/CentOS 4.x, 5.x, 6.x, 7.x, SLES 11, 12, 15, Ubuntu 12, 14, 16, 18, Windows 7, and Windows Server 2003 R2, 2008 R2, 2012 R2, 2016, 2019.
- The VM instances will be forced to shut down during the V2V migration. Therefore, pay attention to the business impact.

Note:

The system firstly attempts to shut down the VM instances gently. If the shutdown fails, the system will perform force shutdown.

- The type of the source primary storage is not enforced. The type of the destination primary storage can be LocalStorage, NFS, Ceph, or SharedBlock.
- For Windows VM instances, the Windows VirtIO driver is automatically installed during the migration. This improves the NIC and disk efficiencies.
- You can perform V2V migration for VM instances booted by UEFI. After the migration, these VM instances are also booted by UEFI.

Source Cloud Platform: KVM

You can migrate VM instances from a KVM platform to the current Cloud by creating a migration task.

- You can perform bulk V2V migrations for VM instances, and customize configurations of the migrated VM instances.
- You can migrate the VM instances that are running or paused. Do not power off the VM instances to be migrated.
- You can perform V2V migrations for VM instances booted by UEFI. After the migration, these VM instances are also booted by UEFI.
- The type of the source primary storage is not enforced. The type of the destination primary storage can be LocalStorage, NFS, Ceph, or SharedBlock.
- For different types of source primary storages or destination primary storages, the libvirt version and QEMU version must meet the following requirements:
  - If either the source primary storage or destination primary storage is Ceph, use libvirt 1.2.16 and QEMU 1.1 or their later versions.
  - If neither the source primary storage nor destination primary storage is Ceph, use libvirt 1.2.9 and QEMU 1.1 or their later versions.

2.2.2.8.2 V2V Conversion Host

Before you can perform V2V migrations, specify a host in a destination cluster as the V2V conversion host.

- A V2V conversion host must have sufficient hardware resources, such as network bandwidth and disk space. The following table lists the minimum configuration requirements.

Table 2-5: Minimum Configuration Requirements for V2V Conversion Host

Hardware	Configuration Requirements
CPU	Minimum 8 cores
Memory	Minimum 16 GB
Network	Minimum 1 Gigabyte NIC
Storage	Minimum 50 GB for the rest of storage spaces **Note:** You can modify the storage configuration according to the number of VM instances to be migrated

- The type of the V2V conversion host must be consistent with that of the source cloud platform.
- You can set an independent migration network and a network QoS for a V2V conversion host to control transmission bottlenecks and to improve migration efficiencies.

2.2.3 Operational Management

2.2.3.1 Tenant Management

Tenant Management allows users to create and manage their organization structures based on their actual business scenarios. It also provides features such as project-based resource access control, ticket management, and independent zone management.

The Tenant Management feature is provided in a separate module. Before you can use this feature, you need to purchase the Plus License of Tenant Management, in addition to the Base License.

Definitions

Definitions related to Tenant Management:

- Personnel and Permissions: The Tenant Management system is structured on the basis of personnel and permissions. You can create departments and roles based on your business needs, and grant a variety of permissions to your users.
- Organization: Organization is the basic unit in Tenant Management. You can create an organization or synchronize an organization through 3rd-party authentication. The organizations can be categorized into the default department and the customized department. You can customize a new team and a sub-department. The new team, usually a company or subcompany (subsidiary), can be used to create multi-level departments. An organizational structure tree is displayed in cascade, and you can directly get a complete picture of the organization structure.

  Note:
  Notice that project members can only view the organization structure where their team belongs to.

- User: A user is a natural person that constructs the most basic unit in Tenant Management. There are local user and the 3rd-party user on NexaVM Cloud.

  — Local User: A user that is created on the Cloud. A local user can be added to an organization or a project, and attached to a role.

  — 3rd-Party User: A user is that is synchronized to the Cloud through 3rd-party authentication. A 3rd-party user can be added to an organization or a project, and attached to a role, and changed to a local user.

  Note:

  - To log in to the Cloud, tenant management users need to use the Tenant login entry.
    - Local users log in to the Cloud via the Local User entry.
    - AD/LDAP users log in to the Cloud via the AD/LDAP User entry.
    - OIDC/OAuth2/CAS users log in to the Cloud from the 3rd-party application without the password.
  - The admin and platform manager can view the list of all users.
  - If you created an organizational structure tree on the Cloud, platform members can view only the list of users belonging to the organizational structure. If you did not create any organizational structure tree, platform members can view all users.

- User Group: A user group is a collection of natural persons or a collection of project members. You can use a user group to grant permissions.
- Role: A role is a collection of permissions that can be granted to users. A user that assumes a role can call API operations based on the permissions specified by the role. Roles are categorized into platform roles and project roles.
  - Platform Role: After a user has a platform role attached, the user will have the management permission of the corresponding zone. Permissions of a platform role take effect only in the zone managed by the user.
  - Project Role: After a user joins a project and have a project role attached, the user will have the permission to use the project and manage the data in the project.

  Note:

  - One user can have both platform roles and project roles attached.
  - One user can have more than one platform role or project role attached.
  - In a project, if a user has multiple project roles attached, the user will have all the permissions attached to the project roles.
- 3rd-party Authentication: The 3rd-party authentication service provided by the Cloud. It supports seamless access to 3rd-party authentication systems. Through the service, related users can directly log in to the Cloud and manage cloud resources. Currently, AD/LDAP/OIDC/OAuth2/CAS servers can be added.
  - AD authentication:

    Active Directory (AD) is a directory service designed for Windows Standard Server, Windows Enterprise Server, and Windows Datacenter Server. AD provides an independent, standard login authentication system for increasingly diverse office applications.

    AD users or organizations can be synchronized to the user list or organization of NexaVM Cloud via an AD server, while specified AD login attributes can be used to directly log in to NexaVM Cloud.
  - LDAP authentication:

    Lightweight Directory Access Protocol (LDAP) can provide a standard directory service that offers an independent, standard login authentication system for increasingly diverse office applications.

LDAP users can be synchronized to the user list of NexaVM Cloud via an LDAP server, while specified LDAP login attributes can be used to directly log in to NexaVM Cloud.

- OIDC authentication:

  OpenID Connect (OIDC) is a set of authentication protocols based on the OAuth2 protocol , and it allows the clients to verify the user identity and obtain basic user configuration information.

  The user information can be synchronized to the Cloud according to the mapping rules via an OIDC server, and users of the OIDC authentication system can log in to the Cloud without the password.

- OAuth2 authentication:

  Open Authorization 2.0 (OAuth2) is a set of authorization protocol standards that can authenticate and authorize users to access related resources. The Cloud currently only supports authorization through the authorization code.

  The user information can be synchronized to the Cloud according to the mapping rules via an OAuth2 server, and users of the OAuth2 authentication system can log in to the Cloud without the password.

- CAS authentication:

  Central Authentication Service (CAS) is a set of single sign-on protocols that allow website applications to authenticate users.

  The user information can be synchronized to the Cloud according to the mapping rules via a CAS server, and users of the CAS authentication system can log in to the Cloud without the password.

- Project Management: Project management allows you to schedule resources based on projects. You can create an independent resource pool for a specific project. By this way, you can better manage the project lifecycle (including determining time, quotas, and permissions) to improve cloud resource utilizations at granular, automatic level and strengthen mutual collaborations between project members.
- Project: A project is a task that needs to be accomplished by specific personnel at a specified time. In Tenant Management, you can plan resources at the project granularity and allocate an independent resource pool to a project. The word **Tenant** in Tenant Management mainly refers to projects. A project is a tenant.

- When you create a project, you need to specify the resource quotas and reclaim policy, and add project members.
- The basic resources (instance offering, image, network, and other resources) on the Cloud are suggested to shared or created in advance.

- Ticket Management: To better provide basic resources efficiently for each project, project members (project admins, project managers, or regular project members) can apply for tickets to obtain cloud resources. Tickets are reviewed and approved according to custom ticket review processes of each project. Finally, the admin, project admins, department managers, and the customized approvers approve the tickets. Currently, five types of ticket are available: apply for VM instances, delete VM instances, modify VM configurations, modify project cycles, and modify project quotas.
- Process Management: Process management is part of ticket management that manages the processes related to the resources of projects. Processes can be categorized into default processes and custom processes.
  - Default process: The project member submits a ticket to the admin, and then the admin approves the ticket. This process applies to the following scenarios:
    - The tickets that are not configured with a ticket process.
    - The tickets which apply for modifications on the project cycle.
    - The tickets which apply for modifications on the project quota.
    - If the custom ticket process is deleted, the tickets will be resubmitted automatically via the default ticket process.
  - Custom process: The project member submits a ticket. The project member makes process settings via process management. Finally, the admin or project admin approves the ticket. This process applies to the following scenarios:
    - The tickets created to apply for VM instances, delete VM instances, and change VM configurations will be prioritized to be submitted via the configured, custom ticket process .
    - If you modify the valid ticket process, the tickets will be automatically resubmitted via this modified, custom ticket process.
    - If you modify the invalid ticket process, you need to resubmit the tickets manually by using this modified, custom ticket process.

- My Approval: In the Cloud, only the administrator and project administrators are granted approval permissions. the administrator and project administrators can approve or reject a ticket. If a ticket is approved, resources are automatically deployed and allocated to the specified project.

Note:
The platform admin and regular platform members do not have the permission for ticket management, and the menu My Approval is not supported for these two roles.

Architecture

The Tenant Management mainly includes four subfeatures, including **project management**, **ticket management**, **independent zone management**, and **3rd-party authentication**.

- **Platform Management**:

  To effectively manage the Cloud, the platform user (platform admin/regular platform member) can cooperate with the super administrator to manage and operate the Cloud together. NexaVM Cloud provides various system roles such as Platform Admin Role and Dashboard Role. You can also satisfy various usage scenarios by creating custom roles at the API level.

- **Project Management**:

  The project management is project-oriented to plan for resources. Specifically, you can create an independent resource pool for a specific project. Project lifecycles can be managed (including determining time, quotas, and permissions) to improve cloud resource utilizations at granular, automatic level and strengthen mutual collaborations between project members.

- **Ticket Management**:

  To better provide basic resources efficiently for each project, project members (project admins , project managers, or regular project members) can submit tickets to obtain cloud resources . Tickets are reviewed and approved according to custom ticket review processes of each project. Finally, the admin, project admins, department managers, and the customized approvers approve the tickets. Currently, five types of ticket are available, including applying for VM instances, deleting VM instances, modifying VM configurations, modifying project cycles, and modifying project quotas.

- **Independent Zone Management**:

  Usually, a zone corresponds to an actual data center in a place. If you isolated resources for zones, you can specify the corresponding zone admins for each zone to achieve independent

managements of various machine rooms. In addition, the admin can inspect and manage all zones.

- **3rd-Party Authentication**:

  The 3rd-party authentication is a third-party authentication service provided by NexaVM Cloud. You are allowed to seamlessly access the third-party login authentication system. The corresponding account system can directly log in to the Cloud to conveniently use cloud resources. Currently, you can add an AD/LDAP/OIDC/OAuth2/CAS server.

Differences in Roles and relevant Permissions

Definitions related to Tenant Management Account System:

- admin: A super administrator who owns all permissions. Usually, the admin is the IT system administrator who have all the permissions.
- Local User: A user that is created on the Cloud. A local user can be added to an organization, added to a project, and attached to a role.
- 3rd-Party User: A user that is synchronized to the Cloud through 3rd-party authentication. A 3rd-party user can be added to an organization, added to a project, and attached to a role.
- Platform User: A user that is not added to a project yet, including platform admin and the regular platform member.
- Platform Admin: A user that has the platform admin role attached. A platform admin who has been allocated a specified zone or all zones manages the data center of the allocated zone or zones.
- Head of Department: The admin can assign a head for the department, and this role is used for identification only. When a head of department becomes a project member, the head of a department has the permission to check department bills.
- Project User: A user who has joined a project, including project admin, project operator, and regular project member.
- Project Admin: A user that has the project admin role attached. A project admin is responsible for managing users in a project, and has the highest permission in a project.
- Project Manager: A user that has the project manager role attached. A project manager assists project admins to manage projects. One or more project members in the same project can be specified to act as project managers.
- Department Manager: The admin can assign a department manager for the new team. It is a type of platform role and is responsible for the operation management of the entire department

, including project management, ticket management, checking bills, and department critical resource monitoring.

- Root Role: The root role is used to limit the permission scope of the custom role. The permission of a custom role is inherited from its root role, and is a subset of the root role permission.
- Quota: A measurement standard that determines the total quantity of resources for a project. A quota mainly includes the VM instance count, CPU count, memory capacity, maximum number of data volumes, and maximum capacity of all volumes.
- Project Reclaim Policy: You need to specify a project reclaim policy when you create a project. There are three types of project reclaim policy, including unlimited, reclaim by specifying time, and reclaim by specifying cost.
  - Unlimited: After you create a project, resources within the project will be in the enabled state by default.
  - Reclaim by Specifying Time:
    - When the expiration date for a project is less than 14 days, the smart operation assistant will prompt you for **The license will be expired** after a project member logs in to the Cloud.
    - After the project expired, resources within the project will be collected according to the specified policy. The policy includes disabling login, preventing project members from logging in to the Cloud, stopping resources, and deleting projects.
  - Reclaim by Specifying Cost: When the project spending reaches the maximum limit, resources within the project will be collected according to the specified policy. The policy includes disabling login, preventing project members from logging in to the Cloud, stopping resources, and deleting projects.
- Access Control: When you create a project, you can specify whether to allow or prohibit project members to or from logging in to the project within a specified time period. There are two types of access control policy: login allowed time and login prohibited time.
  - Login Allowed Time: You can set the time when members in the project can log in to the project by day or week. After setting, the project members can log in to the project only during the login allowed time period.
  - Login Prohibited Time:You can set the time when members in the project cannot log in to the project by day or week. After setting, the project members cannot log in to the project during the login prohibited time period.

- Security group constraint: If you enable security group constraint, when a project member creates a VM instance, the VM instance must have one or more security groups attached.
  - Before you can enable security group constraint for the project, make sure that the project security group quota is set to 1 or higher.
  - If you enable the security group constraint for the project, a default security group is created when the project is created.

The tenant management system grants users a variety of permissions. The permissions of different user roles are as follows:

- **Differences in Accounts Login in Tenant Management**
  - Admin can log in to the Cloud via Account Login.

    By using Chrome or Firefox, go to the Account Login page via *http://management_node_ip:5000/#/login*. To log in to the Cloud, the admin must enter the corresponding user name and password.

    Figure 2-36: Main Login Page

    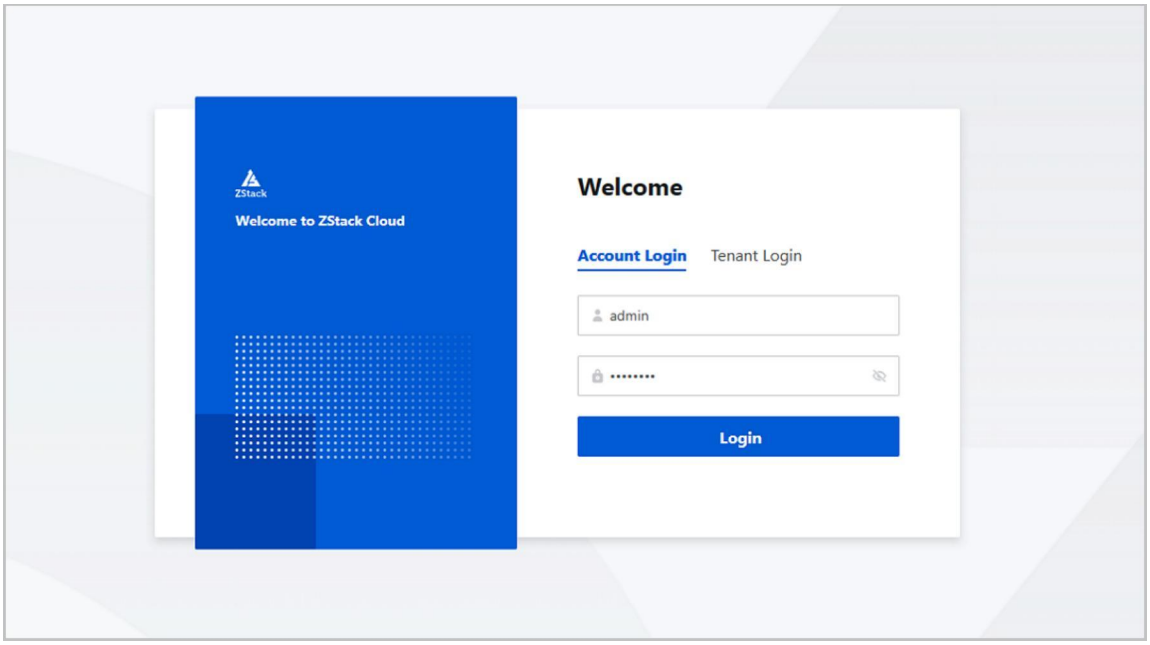

  - For users (platform admin, platform user, project admin, project manager, regular project member, or department manager), log in to the Cloud via Project Login.

    By using Chrome or Firefox, go to the Project Login page via *http://management_node_ip:5000/#/project*. To log in to the Cloud, enter the corresponding user name and password. Specifically, the Cloud has two login entrances for Project Login as follows:

- Local user: the user created on the Cloud. Log in to the Cloud via Local User.
- AD/LDAP user: the 3rd-party user synchronized to the Cloud via the 3rd-party authentication. Log in to the Cloud via AD/LDAP User, as shown in Project Login Page.

After the successful login, you can select the platform or project to be managed to log in to the corresponding management interface.

Figure 2-37: Tenant Login Page

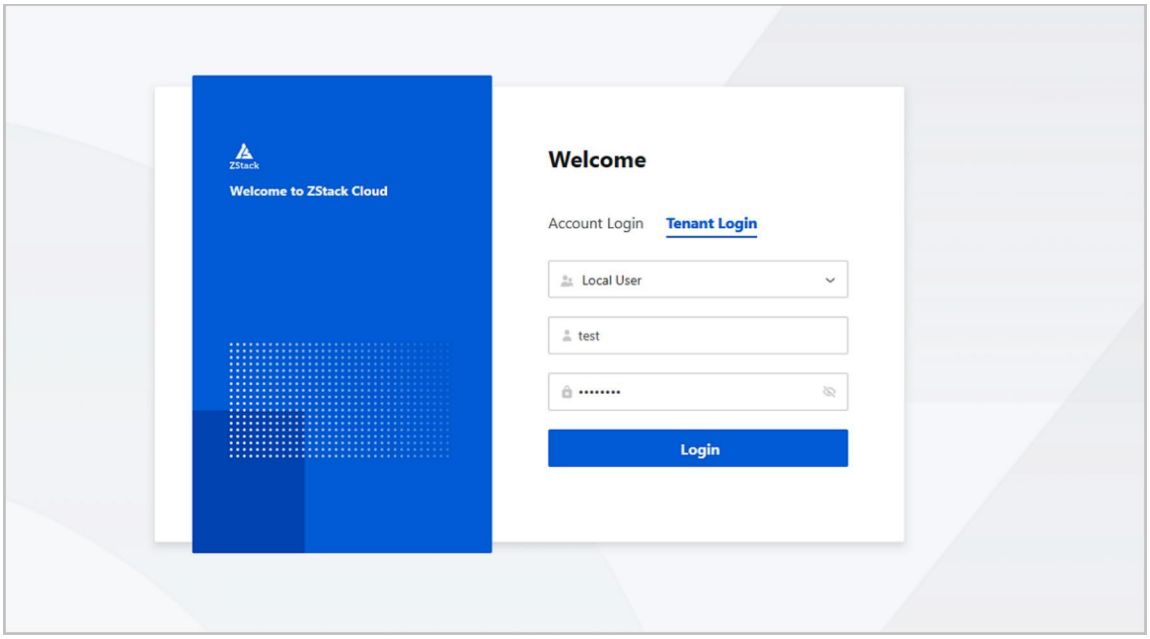

- **Feature Differences from Various Perspectives**

Feature Menu	admin/Platform Admin/Regular Platform Member	Project Admin / Project Manager	Department Manager	Regular Project Member
Organization	○	○	○	○
User	○	○	○	×
Role	○	○	○	○
Project Member	×	○	×	○
User Group	○	○	○	○
3rd-Party Authentication	○	×	×	×
Project	○	×	○	×

Feature Menu	admin/Platform Admin/Regular Platform Member	Project Admin / Project Manager	Department Manager	Regular Project Member
Process Management	○	×	×	×
My Tickets	×	○	×	○
My Approval	○	○	○	×

- **Differences in Permissions of Platform/Project Roles**
  - Platform Roles: admin, platform admin, department manager, and regular platform user. The permissions corresponding to these roles are differentiated as follows:

Role	Difference
admin	A super administrator who owns all permissions.
Platform Admin	A platform admin is a type of administrator who has been allocated a specified zone or all zones, and assists the admin to jointly manage the Cloud. A platform admin has all the permissions that the admin has, except the following: • A platform admin is allocated a specified zone or all zones, and has the permissions to manage resources in the zone or zones only. Currently, a platform admin is not granted relevant permissions to create or delete zones. • A platform admin does not have the permissions related to ticket management, and the menu **My Approval** is not displayed for this role. • A platform admin does not have the permissions related to certificate management, and cannot perform actions such as uploading a certificate.
Department Manager	The department manager is a role who has been allocated a specified department, which can be designated by the admin for the new team and responsible for managing the whole department. A department manager has the following permissions: • View homepage: Allows you to view the summary of project resources in the department under the management only. • View the Cloud monitor: Allows you to view the monitoring information of critical resources of the department under your management.

Role	Difference
	• View organizations: Allows you to view the organizational structure of the Cloud, but not to perform related operations. • View users: Allows you to view the user information on the Cloud, but not to perform related operations. • View user groups: Allows you to view the user group information , but not to perform related operations. • Viewing roles: Allows you to view the system project roles of the Cloud, the project roles whose owner is the admin, and the project roles whose owner is the management department (and sub-departments). • View projects and project-based operations: For projects under the managed department (and sub-departments), you can view, edit, and add project members. Setting a department, changing billing prices, generating project templates, and setting logon time limits for projects are not supported. • Ticket approval: Supports ticket approval, but the menu **Process Management** is not displayed. • View/Export bills: Allows you to view or export project bills and departmental bills of the department (and sub-departments) under your management.
Regular Platform Member	Platform members other than the platform admin. A Platform member has all the permission that the admin has, except the following: • A regular platform member does not have the permissions related to ticket approval, and the menu **My Approval** is not displayed for this role. • A regular platform member can view users who are in the same organizational structure only. • Ungranted permissions.

- Project Roles: project admin, project manager, and project member. The permissions corresponding to these roles are differentiated as follows:

  — A project admin can specify one or more project members in the same project to act as project managers, assisting project admins to manage projects.

  — A project manager has all the permissions that a project admin has, but

Advantages

The Tenant Management of NexaVM Cloud has the following advantages:

- Full-featured: Tenant Management provides users with a range of features such as organizati on structure managements, project-based resource access control, ticket management, and independent zone management.
- User-friendly: Tenant Management allows you to manage the operation permissions of different roles in a multi-level organizational structure, making the organizational management more flexible and user-friendly.
- Cost-effective: Each organization has different kinds of departments. In a traditional IT company, resources are allocated to these departments based on their actual needs, and permissions are assigned as needed as well. Against the backdrop of cloud migration, the management over the departments is achieved on the cloud to minimize the management costs.

Scenarios

Each organization has its own administrative departments. In a traditional IT company, resources are allocated to administrative departments based on their actual needs, and permissions are assigned as needed as well. After companies migrate their business to the cloud, they expect to enjoy the same experience in resources allocation and permissions assignment on the cloud, which is compatible with the management by administrative departments.

The Tenant Management of NexaVM Cloud provides users with a range of features such as organization structure managements, project-based resource access control, ticket management, and independent zone management. Through the division of the organizational structure, it provides the same management as the administrative department and minimizes the management costs.

2.2.3.1.1 Organization

Tenant Management provides an organization management feature for enterprise users, where an organizational structure tree is displayed in cascade and you can directly get a complete picture of the enterprise organization structure. Enterprise Management mainly includes the following concepts:

The concepts of an organization is shown in *Associated Concepts of Organization*.

Figure 2-38: Associated Concepts of Organization

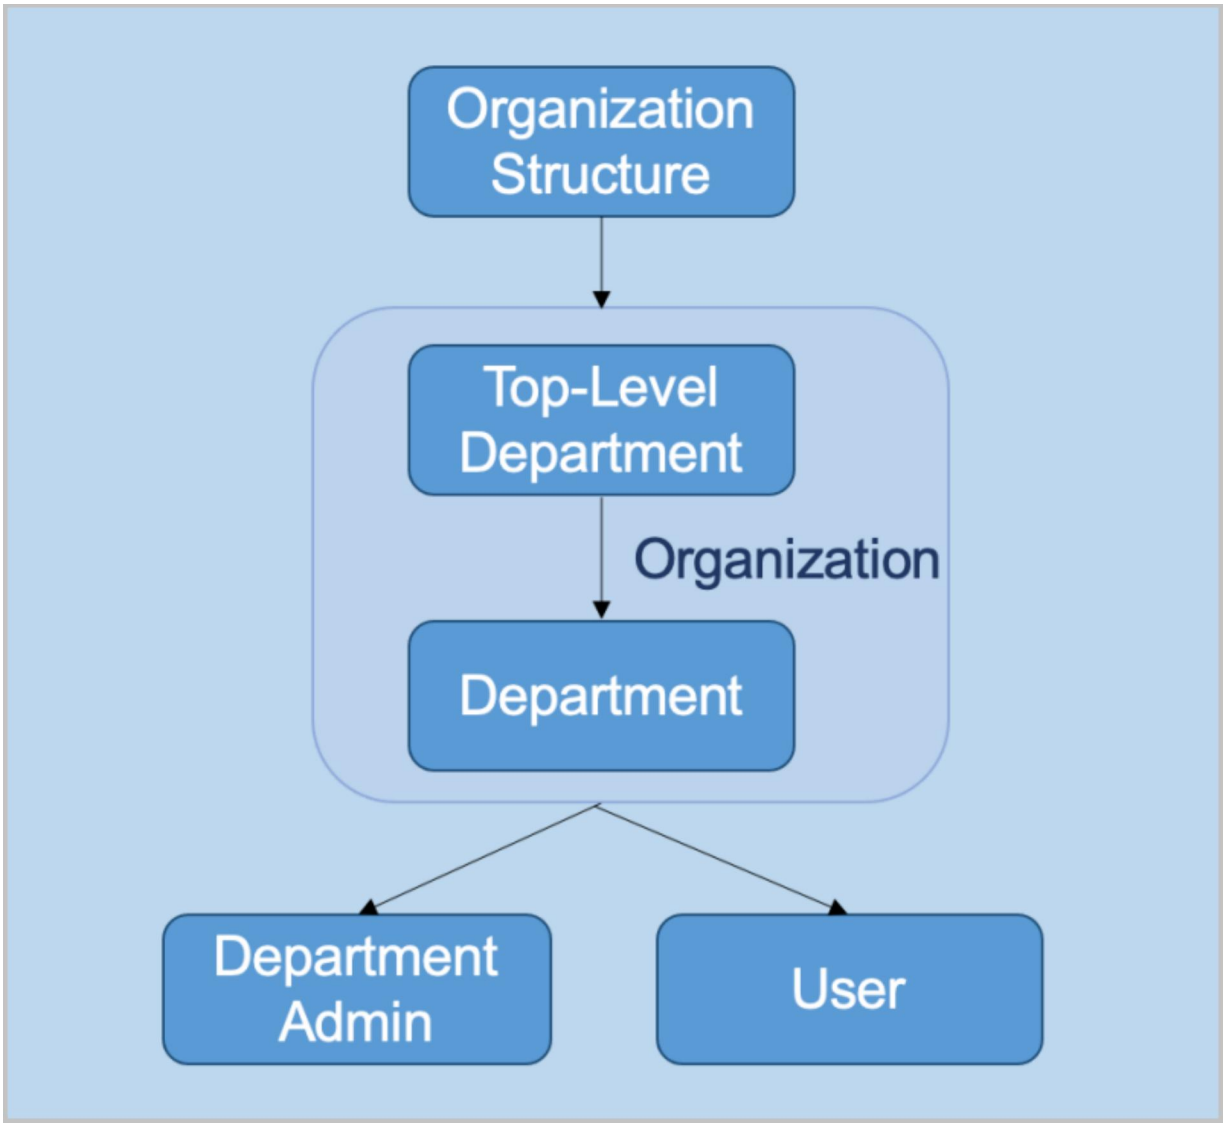

2.2.3.1.2 User

A user is a natural person that constructs the most basic unit in Tenant Management.

Users in NexaVM Cloud can be divided into different types based on where the user is created and whether the user joined a project.

- User type based on where the user is created:
  - Local user: Users created in the Cloud. You can add a local user to an organization or project, or attach a role to a local user.
  - 3rd-party user: Users synchronized to the Cloud through 3rd-party authentication. You can add a 3rd-party user to an organization or project, attach a role to a 3rd-party user, or change a 3rd-party user to local user.

Note:

- To log in to the Cloud, tenant management users need to use the project login entry.
  - Local users log in to the Cloud via the Local User entry.
  - AD/LDAP users log in to the Cloud via the AD/LDAP User entry.
  - OIDC/OAuth2/CAS users log in to the Cloud from the 3rd-party application without the password.
- The admin and platform administrator can view the list of all users.
- If you created an organizational structure tree in the Cloud, platform members can view only the list of users belonging to the organizational structure. If you did not create any organizational structure tree, platform members can view all users.

- User type based on whether the user joined a project:
  - Platform member: A user that is not added to a project yet, including platform manager and the regular platform member.
  - Project member: A user that has joined a project, including project admin, project manager, and regular project member.

2.2.3.1.3 Role

A role is a collection of permissions used for entitling users to manage resources by calling associated APIs. A role has two types, including system role and custom

- Platform role: After a user has a platform role attached, the user will have the management permission of the corresponding zone. Permissions of a platform role take effect only in the zone managed by the user.
- Project role: After a user and its member group join a project and have a project role attached, the user will have the permission to use the project and manage the data in the project.

Note:

- One user can have two types of roles attached.
- One user can have more than one platform role or project role attached.
- In a project, if a user and its member group have multiple project roles attached, the user and its member group will share all the permissions of the user and the member group.

The same user supports binding two role types at the same time.

The same user supports binding multiple platform roles or project roles.

In a project, if a user and member group are bound to multiple project roles, the permissions they have are the full set of all project roles.

2.2.3.1.4 3rd Party Authentication

3rd-Party Authentication is a third-party authentication service provided by NexaVM Cloud. With this service, NexaVM Cloud can seamlessly connect the third-party login authentication system and the corresponding account system can directly log in to the Cloud to conveniently use cloud resources. Currently, you can add an AD/LDAP server.

- AD authentication:

  Active Directory (AD) is a directory service designed for Windows Standard Server, Windows Enterprise Server, and Windows Datacenter Server. AD provides an independent, standard login authentication system for increasingly diverse enterprise office applications.

  AD users or organizations can be synchronized to the user list or organization of NexaVM Cloud via an AD server, while specified AD login attributes can be used to directly log in to NexaVM

  Cloud.

- LDAP authentication:

  Lightweight Directory Access Protocol (LDAP) can provide a standard directory service that offers an independent, standard login authentication system for increasingly diverse enterprise office applications.

  LDAP users can be synchronized to the user list of NexaVM Cloud via an LDAP server, while specified LDAP login attributes can be used to directly log in to NexaVM Cloud.

- OIDC authentication:

  OpenID Connect (OIDC) is a set of authentication protocols based on the OAuth2 protocol, and it allows the clients to verify the user identity and obtain basic user configuration information.

  The user information can be synchronized to the Cloud according to the mapping rules via an OIDC server, and users of the OIDC authentication system can log in to the Cloud without the password.

- OAuth2 authentication:

Open Authorization 2.0 (OAuth2) is a set of authorization protocol standards that can authenticate and authorize users to access related resources. The Cloud currently only supports authorization through the authorization code.

The user information can be synchronized to the Cloud according to the mapping rules via an OAuth2 server, and users of the OAuth2 authentication system can log in to the Cloud without the password.

- CAS authentication:

  Central Authentication Service (CAS) is a set of single sign-on protocols that allow website applications to authenticate users.

  The user information can be synchronized to the Cloud according to the mapping rules via a CAS server, and users of the CAS authentication system can log in to the Cloud without the password.

2.2.3.1.5 Project Management

Tenant Management provides the project management feature for enterprise users.

Project management allows you to schedule resources based on projects. Specifically, you can create an independent resource pool for a specific project. This way, you can better manage the project lifecycle (including determining time, quotas, and permissions) to improve cloud resource utilizations at granular, automatic level and strengthen mutual collaborations between project members.

Concepts of the project management is shown in *Associated Concepts of Project Management*.

Figure 2-39: Concepts of Project Management

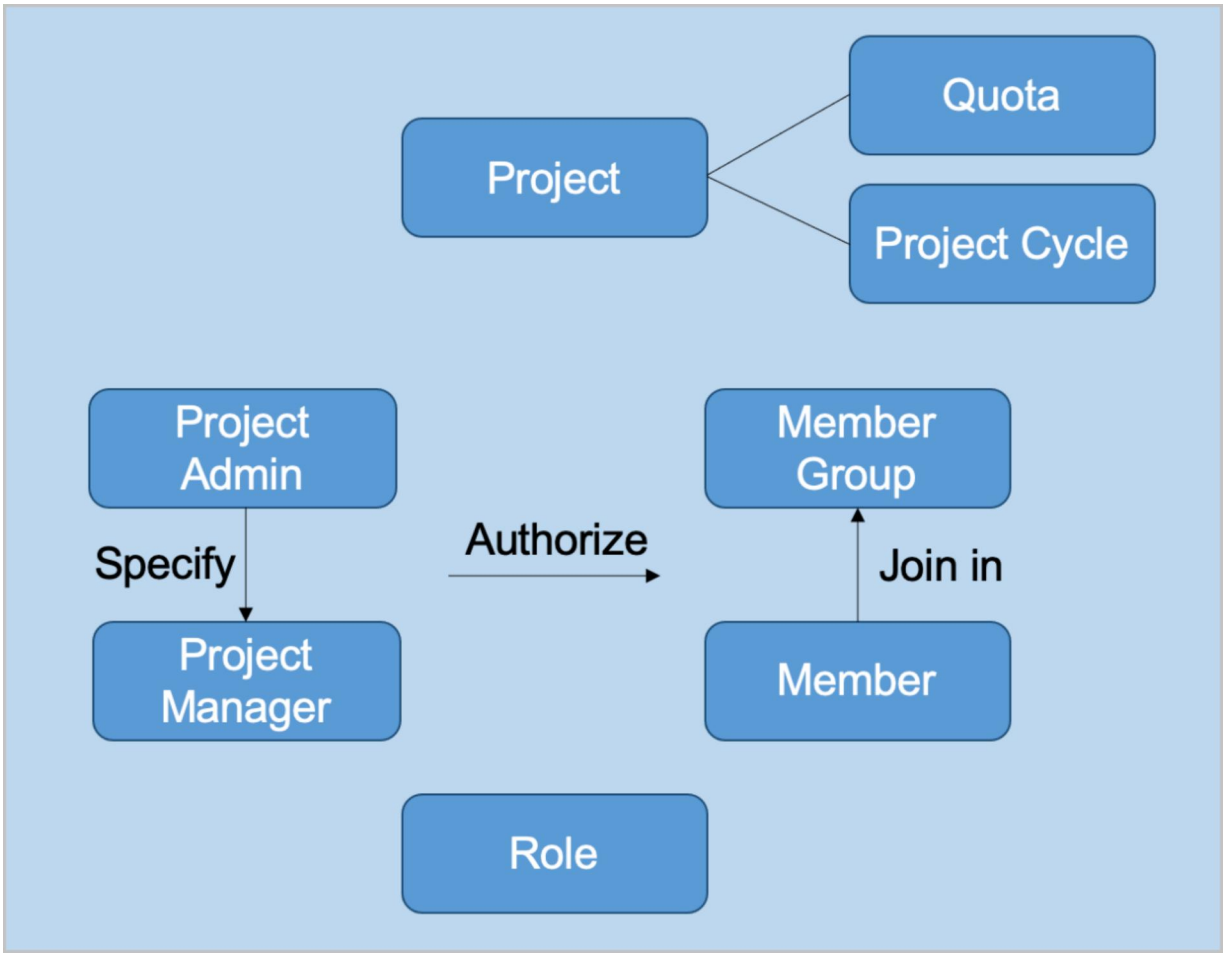

2.2.3.1.6 Ticket Management

The Tenant Management provides the ticket management feature for enterprise users.

To better provide basic resources efficiently for each project, project members (project admins, project managers, or regular project members) can apply for tickets to obtain cloud resources. Tickets are reviewed and approved according to custom ticket review processes of each project. Finally, admins, project admins, or department managers approve the tickets. Currently, five types of ticket are available: apply for VM instances, delete VM instances, modify VM configurations, modify project cycles, and modify project quotas.

The major workflow is shown in *Major Workflow of Ticket Management*.

Figure 2-40: Major Workflow of Ticket Management

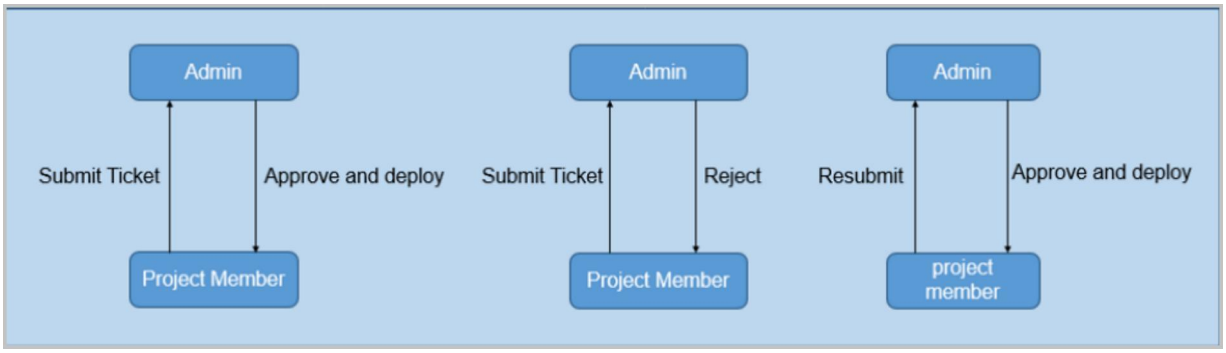

2.2.3.2 Billing Management

2.2.3.2.1 Bills

A bill is the expense of resources totaled at a specified time period. Billing is accurate to the second. Bills can be categorized into project bills, department bills, and account bills.

2.2.3.2.2 Pricing List

NexaVM Cloud provides a quasi-public cloud billing experience. You can customize the unit price for different resources by using a pricing list and obtain related bills after you associate the pricing list with a project or an account. Currently, the following resources in the Cloud can be billed: CPU, memory, root volume, data volume, GPU device, elastic baremetal instances, and public IP (VM IP), and public IP (VIP).

2.2.3.3 Access Control

2.2.3.3.1 Console Proxy

Console proxy allows you to log in to a VM instance by using the IP address of a proxy. You can view the information about the proxy used to launch your VM console.

- The console proxy address only needs to be modified on the management node.
- The address of default proxy is the IP address of the management node.
- You can launch the VM console properly only when the state and status is **Enabled** and **Connected**, respectively.

2.2.3.3.2 Access Key

An AccessKey pair is a security credential that one party authorizes another party to call API operations and access its resources in the Cloud. AccessKey pairs shall be kept confidential.

NexaVM Cloud provides two types of AccessKey: local AccessKey and third-party AccessKey.

- Local AccessKey:

  A local AccessKey pair consists of an AccessKey ID and AccessKey secret. It is a security credential that the Cloud authorizes a third-party user to call API operations and access its cloud resources. AccessKey pairs shall be kept confidential. An AccessKey pair has the full permissions of its creator.

- Third-party AccessKey:

  A third-party AccessKey pair consists of an AccessKey ID and AccessKey secret. It is a security credential that a third-party user authorizes the Cloud to call API operations and access its cloud resources. AccessKey pairs shall be kept confidential.

Note:

- AccessKey is a key factor for the Cloud to perform security authentication on API requests. We recommend that you keep your AccessKey confidential to maintain securities.
- If your AccessKey is at risk of leakage, we recommend that you delete it in time and create a new one.

2.2.3.4 Application Center

Application Center allows you to add third-party applications to the Cloud and then access the applications by using the Cloud. It extends the functionality of the Cloud.

2.2.4 Settings

2.2.4.1 Sub-Account Management

A sub-account can be created by the admin or synced from a third-party authentication system and is managed by the admin. Resources created under a sub-account are managed by the sub-account. You can use a sub-account to create and delete resources under its management and implement fine-grained control over the permissions on resources.

Figure 2-41: Sub-account Management

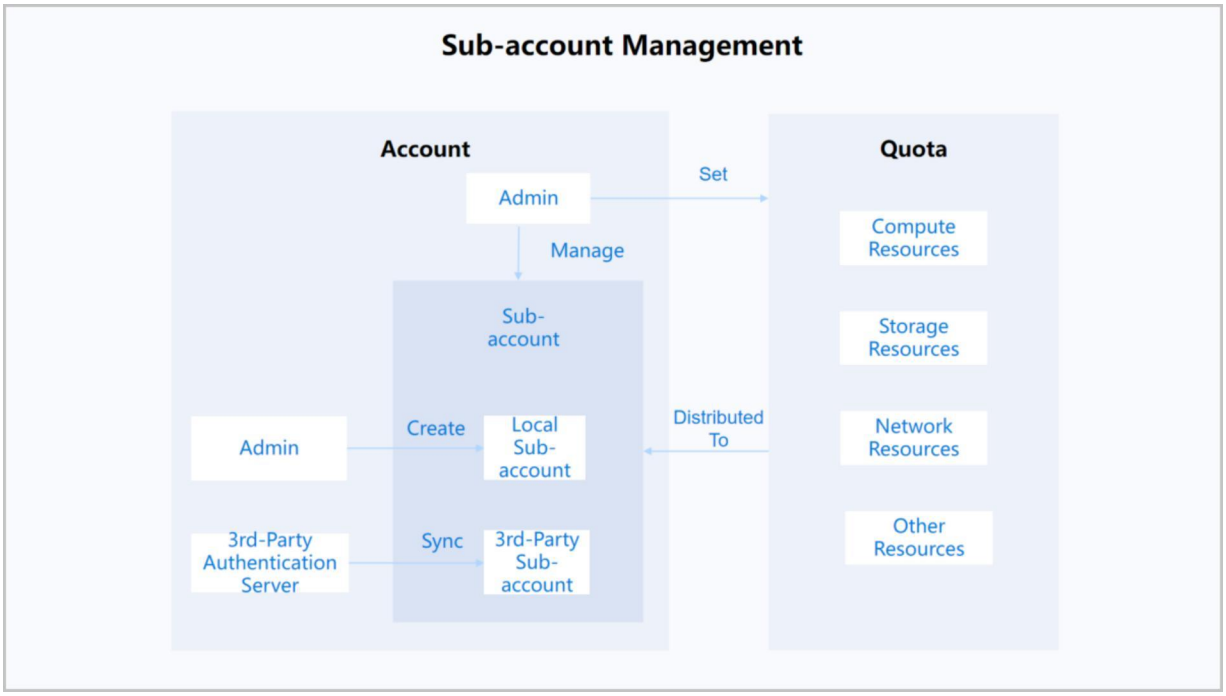

Concepts

- **admin**: The admin has super privileges over resources and shall be owned by the IT system administrator.
  - The admin can share instance offerings, disk offerings, networks, images, and other cloud resources with sub-accounts or revoke the resources from sub-accounts. Sub-accounts can only manage resources to which they are granted access.
  - The admin can modify resource quotas granted to a sub-account based on different business scenarios.
  - After the admin created a VXLAN pool, sub-accounts can create VXLAN networks based on the VXLAN pool.
  - Changing the owner of a VM instance will change the owner properties of the EIPs associated with the VM instance.
- **Sub-account**:
  - Sub-accounts can be categorized into local sub-accounts and third-party sub-accounts:
    - A local sub-account is created by the admin. A third-party sub-account is synced from a third-party authentication server.

- Third-party authentication: The third-party authentication service, powered by the Cloud, supports seamless access to third-party authentication systems. Through the service, related users can directly login to the Cloud and manage cloud resources. Currently, OIDC servers can be added.
  - OIDC server: A third-party authentication server that applies the OIDC protocol . It authenticates and authorizes third-party users to log into the Cloud without password and syncs user information to the Cloud based on the mapping rule.
- A sub-account has management permissions on VM instances, images, volumes, and security groups created under the sub-account. A sub-account can perform read operations on resources shared by the admin, but cannot delete the resources.
- Deleting a sub-account will delete all resources created by the sub-account, such as VM instances, volumes, and images.
- The names of sub-accounts must be unique.
- Resource quotas that the admin shares with a sub-account is displayed on the homepage of the sub-account.
- Before a sub-account can create a VM instance, the admin must share an instance offering, disk offering, network, and other required resources with the sub-account. Otherwise, a VM instance cannot be created.
- A sub-account can use an image that it adds to the Cloud or use an image shared by the admin.

- **Quota**:

  Resource quotas that the admin shares with a sub-account specify the maximum resources that the sub-account can manage, including computing resource quotas, storage resource quotas, network resource quotas, and other resource quotas.

  The admin uses the preceding resource quota settings to manage the maximum resources granted to sub-accounts. If a resource is deleted but not expunged, the resource still occupies storage space of primary storage and volumes.

2.2.4.2 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.

2.2.4.3 Log Server

A log server is used to collect logs of the management node. You can add a log server to the cloud and use the collected logs to locate errors and exceptions. This makes your O&M more efficient.

2.2.4.4 IP Allowlist/Blocklist

An IP blocklist or allowlist identifies and filters IP addresses that access the Cloud. You can create an IP allowlist or blocklist to improve access control of the Cloud.

2.2.4.5 HA Policy

HA Policy is a mechanism that ensures sustained and stable running of the business if VM instances are unexpectedly or scheduled stopped or are errored because of errors occurring to compute, network, or storage resources associated with the VM instances. By enabling this feature, you can customize VM HA policies to ensure your business continuity and stability.

Concepts

The HA Policy feature involves the following key concepts:

- HA mode: Specifies whether to enable auto restart if VM instances are unexpectedly or scheduled stopped or are errored because of errors occurring to compute, network, or storage resources associated with the VM instances. None and NeverStop are supported:
  - None: VM instances scheduled to be stopped or unexpectedly stopped are not auto restarted.
  - NeverStop:
    - VM instances scheduled to be stopped are auto restarted.
    - Unexpectedly stopped VM instances are auto restarted on another host depending on the failover strategy you configure for them.
- VM Failover Strategy: Specifies whether to migrate a VM instance to another host if errors occur to the compute resource, storage resource, or network resource associated with the VM instance.

  The VM failover mechanism inspects the following resource status:

  - Management Network Connectivity Status:
    - Management network connectivity status indicates the status of the network that connects the management node and the host where VM instances reside.

- This status may turn Abnormal if errors occur to the management node or to the management network.

- Storage Network Connectivity Status:
  - Detects the connectivity status of the network that VM instances use to access the primary storage where the root volumes of these VM instances reside.
  - This status may turn Abnormal if errors occur to the primary storage or to the storage network.
- Business NIC Status:
  - Business NIC status may turn Abnormal if errors occur to the host business NIC or the switch port directly connecting to the host business NIC that is associated with the L2 network of VM instances.

Based on the resource status inspection, the Cloud provides the following truth table for configuring VM failover strategies:

Management Network Connectivity Status	**Storage Network Connectivity Status**	**Business NIC Status**	**Fail Over**
Normal	Normal	Abnormal	Yes/No
Normal	Abnormal	Normal	Yes/No
Normal	Abnormal	Abnormal	Yes/No
Abnormal	Normal	Normal	No

Fundamentals

NexaVM Cloud HA Policy has the following mechanisms:

- The Cloud polls the running status of VM instances. If a VM instance is scheduled or unexpectedly stopped, its HA mode is checked. If the HA mode of the VM instance is NeverStop, then the VM instance is restarted on the current host or another host.

Figure 2-42: VM HA Started After Unexpectedly Stopped

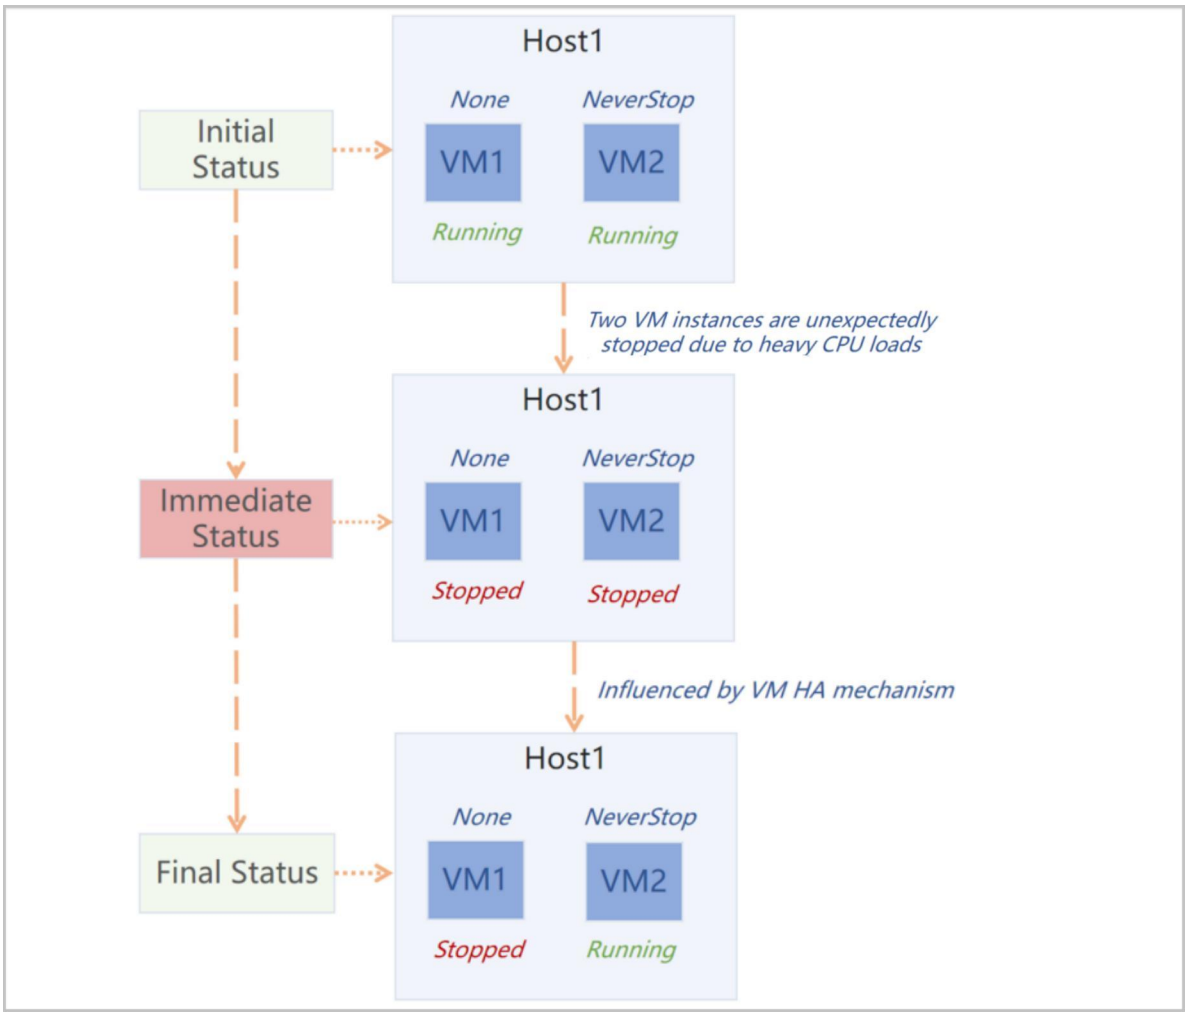

- The Cloud polls the status of the hosts where VM instances reside. Either of the management network connectivity status, storage network connectivity status, and business NIC status of the host turns abnormal, the corresponding VM failover strategy and VM HA mode are checked. If the corresponding failover strategy is Yes and VM HA mode is NeverStop, then related VM instances are migrated to another host.

Figure 2-43: VM HA Started After Host Business NIC Turns Down

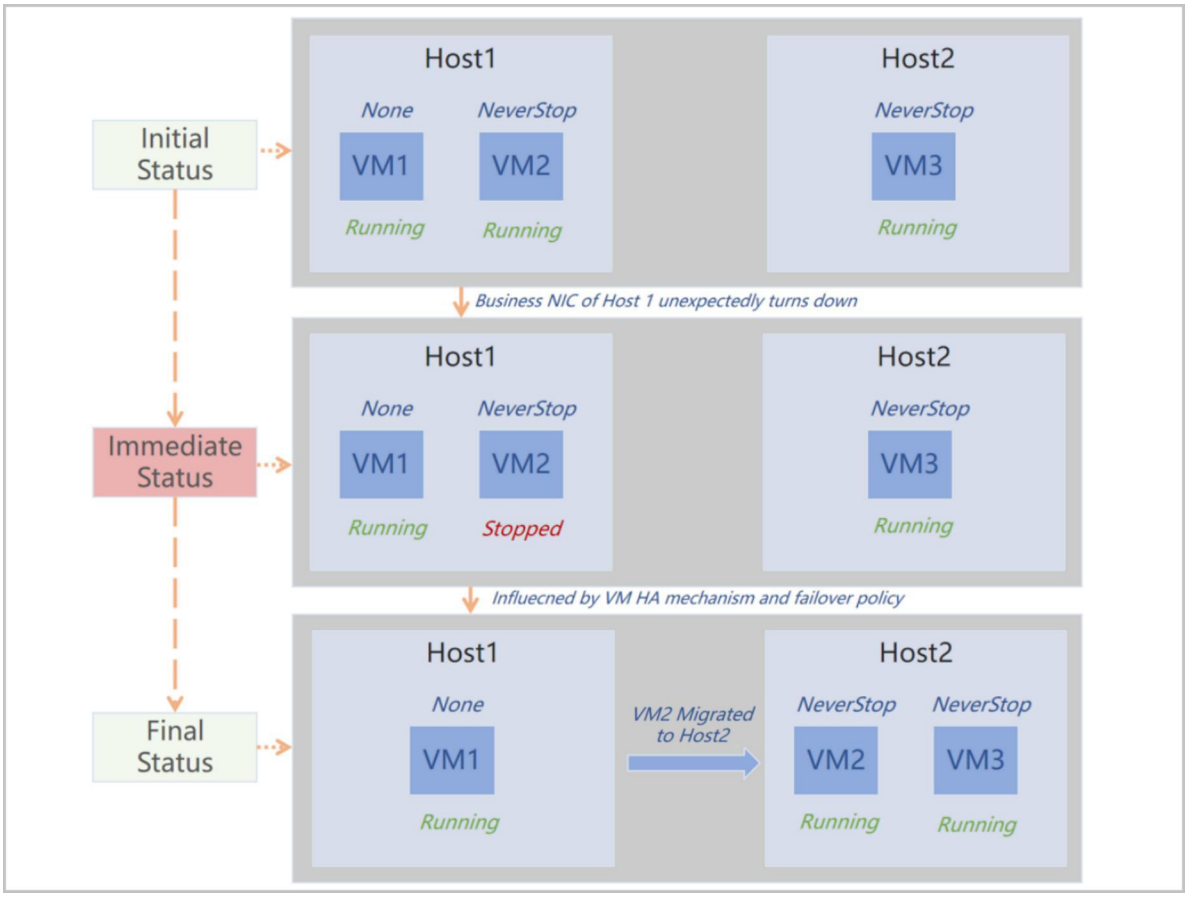

Characteristics

HA Policy has the following characteristics:

- **Comprehensive & Powerful**: Covers all mainstream HA scenarios, including various failures, and ensures the stability and continuity of your business.
- **Flexible & Visualized**: Provides a simple table that allows you to configure VM failover strategies with one click. This table functions together with the HA Mode that can be configured on all and individual VM instances, thus greatly improving the flexibility of your business HA configuration.

Scenarios

The following describes the scenarios of the HA Policy feature.

- Host Business NIC Turns Down:

  If a host business NIC turns down, to ensure high availability of business, all VM instances associated with this NIC are expected to migrate to other hosts.

- For example, your business VM instances are running MySQL database service which is required to achieve high availability. In this case, you can set the HA mode of these VM instances to NeverStop and turn on the switch corresponding to Abnormal Business NIC Status. Then as long as host resources are sufficient, in case that a host business NIC associated with these VM instances turns down, these VM instances will be auto started on other hosts.

- VM Unexpectedly Stops:

  If a VM instance is unexpectedly stopped, it is expected to auto HA start.

  - For example, your VM instances are running important business applications. To ensure business auto-recovery in case of VM stops due to reasons such as host powered-offs or business overloads, you can set the HA mode of these VM instances to NeverStop. Then if these VM instances are stopped, they are auto started.

3 Product Features

Licensing in NexaVM Cloud is supplied in different functionality packages as Base and Plus. This topic describes features covered in the **Enterprise Prepaid** base license and add-on features provided in plus licenses.

For more information about the licensing details, see *License Management*. For differences about features provided in different editions, contact our official sales.

Features in Enterprise Prepaid

Type	Features	Description
Dashboard	Custom Dashboard	Displays multi-dimensional data statistics on cards and allows you to customize your own dashboard by adding and dragging cards.
		Provides a default dashboard for users with different roles.
	Monitor	Displays the platform resources in real time by using monitors with various themes.
		Allows you to switch between the KVM monitor and the vCenter monitor as needed.
		Allows you to switch between zones. You can have the real-time monitoring on all zones or a specific zone.
	API Inspector	Allows you to view the details of API requests that are called by using various methods, including POST, DELETE, PUT, GET, and GET-ZQL, after you perform operations on the UI. NexaVM Cloud supports a browser-based interface using HTML5 or later version for managing and monitoring of server resources.
VM Instance	Bulk Action	Allows you to manage VM instances in bulk.
	Create VM Instance	Allows you to create VM instances through different entries.
		Allows you to specify the root volume capacity and batch attach data volumes via VM creation.
	Import VM Instance	Allows you to import a VM instance on a third-party platform by using the OVF template and customize the configurations of the VM instance.
	View VM Instance	Provides two VM display methods: List View and Directory View.

Type	Features	Description
		Allows you to set a default view for the VM instance page in Global setting or switch view for the current page.
	Manage VM Instance	Allows you to manage the lifecycle of VM instances, such as creating, importing, stopping, booting, rebooting, powering off , recovering, pausing, exporting, and deleting VM instances.
	VM Console	Allows you to access VM instances through terminals without using remote tools.
		Supports three types of console mode: SPICE, VNC, and SPICE+VNC.
		The SPICE protocol supports SSL encrypted channel to further ensure desktop security.
		VNC consoles allow you to paste text.
		Allows you to set the console password, set the console password by force in the Global Setting, and configure the password strategy such as the password complexity and password length in the Global Setting.
	Clone VM Instance without Data Volumes	Copies data in the root volumes of the VM instance only.
		Allows you to clone running, paused, and stopped VM instances on LocalStorage, NFS, SMP, Ceph, and SharedBlock primary storage.
		Supports ImageStore and Ceph backup storage.
		Allows you to choose clone method as needed, including full clone, instant full clone, and linked clone.
		Allows you to set a storage allocation policy, including system allocation and manual allocation.
	Clone VM Instance with Data Volumes	Copies data in the root volumes as well as data volumes of a VM instance.
		Allows you to clone running, paused, and stopped VM instances on LocalStorage, NFS, SMP, Ceph, and SharedBlock primary storage.
		Supports ImageStore and Ceph backup storage.
		Does not clone shared volumes (if any) with VM instances.
		Allows you to choose clone method as needed, including full clone, instant full clone, and linked clone.

Type	Features	Description
		Allows you to set a storage allocation policy, including system allocation and manual allocation.
	Flatten	Allows you to merge snapshots of a VM instance into one flat snapshot to improve resource performance and data security .
		Allows you to unlink the dependency between linked clone VM instances and source VM instances by flattening to achieve data independence.
	Custom Tag	Allows you to customize tags for VM instances so that you can locate them quickly.
	Change Group	Allows you to create groups to categorize and manage VM instances. You can create up to 4-level groups with the root directory as the first-level group.
		Allows you to manage the lifecycle of a group, such as creating and deleting a group.
		Allows you to specify a group for a VM instance or change the group it belongs to.
	Change Host	Allows you to migrate a VM instance from a host to another without changing the primary storage.
		Supports hot migration and cold migration.
		Hot migration: Migrates a VM instance in the running state. Hot migration applies to all types of the primary storage.
		Allows you to hot migrate a VM instance with a vDPA NIC attached if the VM instance is on a LocalStorage or shared primary storage.Re
		If the migration is blocked because the VM instance has high I/O operations for a long time, you can enable auto converge to ensure a smooth migration.
		Cold migration: Migrates a VM instance in the stopped state. Cold migration applies to LocalStorage primary storage only.
		Allows you to cold migrate a VM instance with a vDPA NIC attached if the VM instance is on a LocalStorage primary storage.
		Allows you to cold or hot migrate a VM instance based on the workloads of the destination host.

Type	Features	Description
	Change Primary Storage	Allows you to migrate a VM instance from a primary storage to another without changing the host.
		Allows you to migrates valid data, and the migrated VM instance follows the provisioning type of the target primary storage.
		Supports hot migration and cold migration across SharedBlock primary storage.
		Supports hot migration across SharedBlock and Ceph primary storage.
		Hot migration: Migrates a VM instance in the running state.
		Snapshots of the VM instance to be migrated will not be saved after the hot migration across SharedBlock primary storage or across SharedBlock and Ceph primary storage.
		If you hot migrate a VM instance from a SharedBlock primary storage to a Ceph primary storage, you can specify a root volume pool or data volume pool for the volumes to be migrated.
		Cold migration: Migrates a VM instance in the stopped state.
		Allows you to hot or cold migrate a VM instance with all attached volumes (excluding shared volumes).
		Allows you to set a limit on the VM storage migration speed.
		Allows you to manually cancel an ongoing migration task.
	Change Host and Primary Storage	Allows you to migrate a VM instance from a host and primary storage to another host and primary storage.
		Supports hot migration and cold migration.
		Hot migration: Migrates a VM instance in the running state.
		Supports hot migration across the same type of primary storage, including Ceph↔Ceph, NFS↔NFS, and SharedBlock↔SharedBlock.
		Snapshots of the VM instances to be migrated will not be saved after the hot migration across the same type of primary storage.
		If you hot migrate a VM instance across Ceph primary storage, you can specify a root volume pool or data volume pool for the volumes to be migrated.

Type	Features	Description
		Allows you to hot migrate a VM instance across different types of primary storage, including Ceph↔SharedBlock, LocalStorage↔SharedBlock, LocalStorage↔Ceph, LocalStorage↔NFS, SharedBlock↔NFS, and Ceph↔NFS.
		Snapshots of the VM instances to be migrated will not be saved after the hot migration across different types of primary storage.
		If you hot migrate a VM instance from a SharedBlock, LocalStorage, or an NFS primary storage to a Ceph primary storage, you can specify a root volume pool or data volume pool for the volumes to be migrated.
		Supports hot migration across Ceph pools within the same Ceph primary storage and allows you to migrate only root volume or migrate data volumes with VM instances.
		Allows you to enable or disable auto-convergence policy during storage migration.
		Allows you to manually specify destination hosts.
		Cold migration: Migrates a VM instance in the stopped state.
		Allows you to cold migrate a VM instance across the same type of primary storage, including Ceph↔Ceph and NFS↔NFS.
		Allows you to cold migrate a VM instance across the same type of primary storage without data volumes.
		If you cold migrate a VM instance across Ceph primary storage, you can specify a root volume pool for the volumes to be migrated.
		Supports cold migration of VM instances (with data volumes) across Ceph pools within the same Ceph primary storage.
		Allows you to clean up raw data after migration to release more space after you confirm the data integrity.
		Allows you to set a limit on the VM storage migration speed.
		Allows you to manually cancel an ongoing migration task.
	Modify Instance Offering	Allows you to modify the instance offering (CPU and memory) of a running or stopped VM instance.

Type	Features	Description
	Set GPU Specification	Allows you to set the GPU specification for a stopped VM instance. You can attach, modify, and detach a physical GPU specification or virtual GPU specification as needed.
	Resize Root Volume	Allows you to expand the root volume of a running or stopped VM instance. The new size takes effect immediately.
	Resize Data Volume	Allows you to expand the data volume of a running or stopped VM instance. The new size takes effect immediately.
	Change Owner	Allows you to change the owner of a running or stopped VM instance.
	Change System	Allows you to change the operating system of a stopped VM instance.
	Reimage VM Instance	Allows you to restore a VM instance to the initial state of the VM image. All the data in the root volume will be overwritten.
	Set Boot Order	Allows you to set the boot order for a VM instance.
		Supports three boot types: boot from hard disk, boot from CD ROM, and boot from network.
	Boot from Host	Allows you to specify a host on which a VM instance boots.
	VM High Availability	Allows you to set VM high availability (HA) so that the VM instance can reboot automatically in case of host exception. You can view the reboot progress on the UI. You can set HA for a VM instance through two entries: Global Setting and VM Setting. The setting takes effect with the following priority: Global Setting < VM Setting.
	Time Synchronization	Allows you to set whether the base time of a VM instance is the same as that of the host.
	SSH Key Attachment/ Detachment	Allows you to attach/detach SSH keys to/from VM instances with the Linux or BSD operating system.
		Allows you to create or delete an SSH key.
	Change VM Password	Allows you to change the password of a Windows or Linux running VM instance.
	Set Hostname	Allows you to set the hostname when you create a VM instance.
	Resource Priority	Allows you to set resource priority (**Normal** and **High**). When resource contention occurs, VM instances with **High**

Type	Features	Description
		resource priority can compete for more resources than those with **Normal** resource priority.
	Bind with Cluster	Allows you to set whether to bind a VM instance with the cluster it is residing on.
		If set to false, the VM instance is not bind with the cluster and allows cross-cluster HA migrations, automatic cross-cluster migrations triggered by the host maintenance mode, as well as manual cross-cluster migrations triggered by actions like " Changing Host" for the VM instance.
		If set to true, the VM cross-cluster migrations are restricted . You can further define which migration operations are restricted through the cluster setting Resource Binding Policy .
		Soft Binding: this policy makes the VM instance be migrated only among the current cluster if the migration is triggered automatically by the VM HA or by a host maintenance mode. Migrations that are triggered manually are not limited by the cluster scope, such as manual VM migrations (Change Host).
		Hard Binding: this policy makes the VM instance always run in the current cluster, and all cross-cluster migrations are denied, including the migrations triggered automatically and manually.
	USB Redirection	Allows you to redirect a USB device from a VDI client to a VDI VM instance.
	VM Snapshot	Allows you to schedule snapshot creation at specified time points to record the state of the root volume, data volume, or memory of an instance before you perform a business-sensitive operation. This allows rollback in case of breakdowns.
		Supports two snapshot types: Single snapshot and snapshot group. The snapshot group allows you to restore a group of VM instances in bulk.
		Allows you to create snapshots for VM instances that are in the running state.
		Allows you to create snapshots for VM instances that are in the stopped state.
		Supports VM auto boot after restoring from snapshots.

Type	Features	Description
		Allows you to delete VM snapshots in bulk.
		Allows you to create a VM instance from a single snapshot or create a VM instance with data volumes from a snapshot group.
	VM Backup	Allows you to create a backup for a running VM instance.
		Allows you to create either an incremental backup or a full backup for a VM instance.
		Allows you to create a backup for a VM instance with its volumes (excluding shared volumes) when the VM instance is in the running state.
		This feature is provided by the Backup Service module.
	VM Image	Allows you to create a template image based on a VM instance so that you can create VM instances in bulk in a custom way.
		Allows you to create a VM image when the VM instance is in the running or stopped state. Supported backup storage: ImageStore and Ceph.
	ISO-based Deployment	Creates VM instances based on an ISO disk which guides the VM system installation.
		Supports multiple ISOs per VM instance, improving the business deployment efficiency.
	Template-based Deployment	Creates VM instances based on a system template.
	Add/Remove VM Scheduling Group	Allows you add a running or stopped VM instance to or remove a VM instance from a VM scheduling group so as to associate with/disassociate from related VM scheduling policies. This way, you can manage the distribution of VM on hosts and ensure high performance and high availability.
	Attach/Detach Volume	Allows you to attach/detach a data volume to/from a running or stopped VM instance. Allows you to optimize drive models and identify a volume by its SCSI WWN.
	Shared Volume	Allows VM instances in Ceph or SharedBlock primary storage to share the same data volume.
	Create Volume Image	Allows you to create an image for root volumes or data volumes attached to a VM instance when the VM instance is in the running or stopped state.

Type	Features	Description
		Before you can create an image for shared volumes on a SharedBlock primary storage, stop all the VM instances attached by the shared volume first.
	Set Volume QoS	Allows you to set QoS for root volumes and data volumes attached by a VM instance when the VM instance is in the running or stopped state.
	Enable/Disable NIC	Allows you to enable or disable NICs of the vNIC type.
	Attach/Detach NIC	Allows you to attach/detach a NIC to/from a running or stopped VM instance. You can set a default NIC.
	Set NIC Type	Allows you to set the NIC type when the VM instance is in the stopped state.
		Allows you to change the NIC type from a VF NIC to a vNIC only.
	Set NIC Model	Allows you to set the NIC model for a running or stopped VM instance. Supported VM NIC models: virtio, rtl8139, and e1000.
		This operation applies to Linux and Paravirtualization operating systems only.
	Change L3 Network for VM NIC	Allows you to change the L3 network for a VM NIC without affecting the hardware information such as the MAC address and PCI address of the NIC.
	Set Network QoS	Allows you to set the network QoS for a running or stopped VM instance.
	Sync NIC Configurations	Allows you to update NIC configurations according to the NIC parameters you set on the Cloud, including IP address, netmask, gateway, DNS, and MTU.
	Read NIC IP Configured in the VM Instance	Allows you to read a NIC IP address configured in the VM instance and make it displayed on and managed by the Cloud. Make sure that the NIC is belong to an L3 network disabled with IP address management and does not have an IP address assigned on the Cloud.
	Customize MAC Address	Allows you to customize a MAC address when you create a VM instance.
		Allows you to change the MAC address when the VM instance is in the stopped state.

Type	Features	Description
	Customize IP Address	Allows you to customize an IP address when you create a VM instance.
		Allows you to change the IP address when the VM instance is in the stopped or running state.
	Attach/Detach EIP	Allows you to attach an EIP to or detach an EIP from a VM NIC.
	VM Multi-Gateway	Allows you to enable multi-gateway by running `NexaVM-ctl`. If enabled, each VM NIC has an independent gateway.
	Create/Delete vDrive	Allows you to create/delete a vDrive for a stopped VM instance. You can attach/detach an ISO to/from a drive.
	Attach/Detach Peripheral Device	Allows you to attach/detach a LUN to/from a running or stopped VM instance.
		Allows you to attach/detach a physical GPU device to/from a running or stopped VM instance.
		Allows you to attach/detach a virtual GPU device to/from a running or stopped VM instance.
		Allows you to attach/detach a USB device to/from a running or stopped VM instance.
		Allows you to attach/detach other peripheral devices, such as Moxa cards, to/from a running or stopped VM instance.
	CPU Model	Allows you to set the CPU model for a VM instance through three entries: Global Setting, Cluster Setting, and VM Setting . The setting takes effect with the following priority: Global Setting < Cluster Setting < VM Setting.
	CPU Pinning	Assigns the virtual CPUs (vCPUs) of a VM instance to specific host pCPUs, which improves VM performance.
	vNUMA Configuration	Allows you to configure vNUMA for a VM instance to generate a topology of virtual NUMA nodes for the VM instance. This topology enables a vCPU on a vNUMA node to primarily access the local memory and thus improves VM performance.
	EmulatorPin Configuration	Allows you to configure EmulatorPin for a VM instance so that all other threads than virtual CPU (vCPU) threads and IO threads of a VM instance are assigned to physical CPUs (pCPUs) of the host.

Type	Features	Description
	VM Performance Optimization	Allows you to install performance optimization tools (GuestTools) for the Qemu Guest Agent installation and internal monitoring of Linux VM instances.
		Allows you to install performance optimization tools (GuestTools) for Windows and Windows Virtio VM instances for Qemu Guest Agent installation and internal monitoring. You can install the Virtio driver with one click to improve the disk and NIC performances.
	Import User Data	Allows you to import user data when you create a VM instance. You can upload user-defined parameters or scripts to customize configurations for VM instances or to accomplish specific tasks.
	BIOS Mode	Inherits the BIOS mode from the image you selected when you create a VM instance. The BIOS mode includes Legacy and UEFI.
		Inherits the BIOS mode of the original VM instance when you create a VM image or clone a VM instance.
		Allows you to change the BIOS mode when the VM instance is in the running or stopped state.
	VM RDP	After RDP is enabled, you can launch the VM console in RDP mode by default in VDI scenarios.
	Anti-Spoofing Mode	Provides IP/MAC anti-spoofing and ARP anti-spoofing. If enabled, VM instances can only communicate with outside networks using the IP/MAC addresses allocated by the Cloud .
	VM Monitoring	External monitoring: Collects the VM data such as CPU, memory, disk I/O, NIC data from hosts by using libvirt.
		Internal monitoring: Collects the VM data such as CPU, memory, and disk size data from VM instances by using an agent. An agent is required for internal monitoring.
	Advanced Settings	Allows you to enable Instance Offering Online Modification for a single VM instance so that you can online modify the instance offering (CPU and memory) for the VM instance.
		Allows you to enable Hyper-V for a Windows VM instance.

Type	Features	Description
		Allows you to disable the hypervisor for a VM instance, to make certain applications skip their virtualization detection on this VM instance.
		Allows you to disable hypervclock for a Windows VM instance.
		Allows you to set the number of queues when VirtIO NIC traffics are allocated to multiple CPUs, which improves the NIC performance.
		Allows you to enable memory reclaim for a VM instance. It monitors in real time the memory usage of VM instances and the host. Its dynamic reclaim and allocation mechanism makes sure the efficient use of host memory
		Allows you to specify the reserve size of memory after you enable memory reclaim.
	Audit	Audits all of the actions performed on VM instances, which effectively ensures the security of the Cloud environment.
	Custom Column	Allows you to customize the items to be displayed on a VM list.
	Export CSV File	Allows you to export the VM information as a CSV table, which helps in statistical analysis and problem diagnosis.
	Resource Deletion Policy	Provides three deletion policies to lower risks caused by misoperations. The policies include Direct, Delay (default), and Never.
		Displays warnings of the consequences on the UI and asks for confirmation before the deletion is completed.
Volume	Bulk Action	Allows you to manage volumes in bulk.
	Create Volume	Provides multiple strategies to create volumes.
	Manage Volume	Allows you to manage the lifecycle of volumes, such as creating, enabling, disabling, and deleting volumes.
	Attach/Detach Instance	Allows you to attach/detach a volume to/from an instance.
	Change Host	Allows you to migrate a volume to another host. This action applies to local primary storage only.
		Allows you to migrate a volume based on the workloads of the destination host.

Type	Features	Description
	Change Primary Storage	Allows you to migrate a volume to another primary storage.
		Allows you to migrate valid data, and the migrated volume follows the provisioning type of the target primary storage.
		Supports volume migration across the same type of primary storage, including Ceph↔Ceph, NFS↔NFS, and SharedBlock↔SharedBlock.
		Allows you to migrate volumes not attached to any instances between Ceph↔Ceph, NFS↔NFS, and SharedBlock↔SharedBlock.
		Allows you to migrate volumes attached to a VM instance in the stopped state across SharedBlock primary storage.
		Supports volume migration across Ceph pools within the same Ceph primary storage.
		Allows you to clean up raw data after migration to release more space after you confirm the data integrity.
	Change Owner	Allows you to change the owner of a volume.
	Resize Volume	Allows you to expand a volume that is not attached to any instance.
		Allows you to expand the volume of a running or stopped instance.
		In Ceph primary storage, allows you to expand a shared volume that is not attached to any instance or is attached to a stopped instance.
	Custom Tag	Allows you to customize tags for volumes so that you can locate them quickly.
	Volume Backup	Allows you to create a backup for a volume that is attached to a running instance.
		Allows you to create either an incremental backup or a full backup for a volume.
		This feature is provided by the Backup Service module.
	Volume Image	Allows you to create a template image based on a volume, and helps you to create volumes in bulk in a custom way.
		Allows you to create an image for a volume that is not attached to any instance.

Type	Features	Description
		Allows you to create an image for a volume that is attached to a running or stopped instance.
		In SharedBlock primary storage, allows you to create an image for a shared volume that is not attached to any instance or is attached to a stopped instance.
		In Ceph primary storage, allows you to create an image for a shared volume that is not attached to any instance or is attached to a running or stopped instance.
	Volume Snapshot	Allows you to schedule snapshot creation at specified time points to record the state of a root volume or data volume before you perform a business-sensitive operation. This allows rollback in case of breakdowns.
		Allows you to restore a volume snapshot as needed.
		Allows you to delete volume snapshots in bulk.
	Set Volume QoS	Allows you to set QoS for volumes.
	Flatten	Allows you to merge snapshots of a volume into one flat snapshot to improve resource performance and data security.
		Allows you to unlink the dependency between linked clone volumes and source volumes by flattening to achieve data independence.
	Shared Volume	Allows you to create shared volumes in Ceph or SharedBlock primary storage.
	Audit	Audits all of the actions performed on volumes, which effectively ensures the security of the Cloud environment.
	Export CSV File	Allows you to export the volume information as a CSV table, which helps in statistical analysis and problem diagnosis.
	Resource Deletion Policy	Provides three deletion policies to lower risks caused by misoperations. The policies include Direct, Delay (default), and Never.
		Displays warnings of the consequences on the UI and asks for confirmation before the deletion is completed.
Image	Bulk Action	Allows you to manage images in bulk.
	Add Image	Allows you to add two types of images: system image (ISO/Image) and volume image (Image).

Type	Features	Description
		Allows you to set the CPU architecture of an image, including x86_64, aarch64, and mips64el. Creating VM instances, creating VM images, and cloning VM instances will inherit the CPU architecture of the original image.
		Allows you to set the image platform, such as Linux,Windows , and Other.
		Allows you to upload an image by using an URL or local browser.
		Allows you to set the BIOS mode for an image, including Legacy and UEFI. Creating VM instances, creating VM images, and cloning VM instances will inherit the BIOS mode of the original image.
	Manage Image	Allows you to manage the lifecycle of images, such as adding , enabling, disabling, and deleting images.
	Change Backup Storage	Allows you to migrate an image to another backup storage. This action applies to Ceph backup storage only.
		Allows you to clean up raw data after migration to release more space after you confirm the data integrity.
	Export Image	Allows you to export an image from an ImageStore or Ceph backup storage.
		Provides the MD5 value of the downloaded image to check the image integrity.
	Sync Image	Allows you to synchronize images among different ImageStore backup storage in the same management node.
	Set Sharing Mode	Allows you to set the sharing mode of an image, including share globally, share to specified projects or accounts, and not share.
	Audit	Audits all of the actions performed on images, which effectively ensures the security of the Cloud environment.
	Resource Deletion Policy	Provides three deletion policies to lower risks caused by misoperations. The policies include Direct, Delay (default), and Never.
		Displays warnings of the consequences on the UI and asks for confirmation before the deletion is completed.
Instance Offering	Bulk Action	Allows you to manage instance offerings in bulk.

Type	Features	Description
	Create Instance Offering	Allows you to select a host allocation policy, including host with minimum number of running VMs (default policy), host with minimum CPU utilization, host with minimum memory utilization, host with maximum number of running VMs, host where the VM is located last time, and random host allocation to create VM instances.
		When the host allocation strategy is host with minimum CPU utilization or host with minimum memory utilization, you can select the mandatory strategy mode or non-mandatory strategy mode (default mode).
		If the Cloud can obtain the host load information, it will create VM instances according to the host allocation strategy. If the Cloud could not obtain the host load information, it will create VM instances according to the strategy mode.
		Allows you to set disk QoS and network QoS for an instance offering.
		Allows you to set advanced parameters through JSON to customize an instance offering.
	Manage Instance Offering	Allows you to manage the lifecycle of images, such as creating, enabling, disabling, and deleting instance offerings.
	Set Sharing Mode	Allows you to set the sharing mode of an instance offering, including share globally, share to specified projects or accounts, and not share.
	Audit	Audits all of the actions performed on instance offerings, which effectively ensures the security of the Cloud environment.
Disk Offering	Bulk Action	Allows you to manage disk offerings in bulk.
	Create Disk Offering	Allows you to set the disk QoS for a disk offering.
		Allows you to set advanced parameters through JSON to customize a disk offering.
	Manage Disk Offering	Allows you to manage the lifecycle of disk offerings, such as creating, enabling, disabling, and deleting disk offerings.
	Set Sharing Mode	Allows you to set the sharing mode of a disk offering, including share globally, share to specified projects or accounts, and not share.

Type	Features	Description
	Audit	Audits all of the actions performed on disk offerings, which effectively ensures the security of the Cloud environment.
GPU Specification	Bulk Action	Allows you to manage physical GPU specifications in bulk. vGPU specifications do not support bulk actions.
	Manage Physical GPU Specification	Automatically detects available physical GPU specifications on the Cloud and lists them in the UI.
		Allows you to enable or disable a physical GPU specification.
	Manage Virtual GPU Specification	Generates virtual GPU specifications from the detected physical GPU specifications and lists them in the UI.
		Allows you to enable or disable a virtual GPU specification.
	Set ROM	Allows you to set ROM for physical GPU specifications for physical GPU passthrough.
	Set Sharing Mode	Allows you to set the sharing mode of a GPU specification, including share globally, share to specified projects or accounts, and not share.
	Audit	Audits all of the actions performed on GPU specifications, which effectively ensures the security of the Cloud environment.
Auto-Scaling Group	Create Auto-Scaling Group	Allows you to set a health check mechanism, including load balancer health check and VM health check to trigger elastic self-healing.
		Allows you to set the resource monitoring and alarm mechanism to trigger elastic scaling. The mechanism includes trigger metrics, scale-out policy, scale-in policy, and whether to enable alarm notification (if enabled, an endpoint is required). The trigger metrics include both external monitoring items (VM Memory Utilization Average, VM CPU Utilization Average) and internal monitoring items (VM Memory Utilization Average, VM CPU Utilization Average) of VM instances. Note that an agent is required for internal monitoring.
		Allows you to set a periodic policy (scale-out policy or scale-in policy) for an auto-scaling group. The scale-out/scale-in cycle can be accurate to minutes with a minimum cycle interval of 15 minutes.

Type	Features	Description
	Manage Auto-Scaling Group	Allows you to manage the lifecycle of auto-scaling groups , such as creating, enabling, disabling, and deleting auto-scaling groups.
	Add/Remove VM Scheduling Group	Allows you add an auto-scaling group to or remove an auto-scaling group from a VM scheduling group so as to associate with/disassociate from related VM scheduling policies. This way, you can manage the distribution of VM instances on hosts and ensure high performance and high availability.
	Change Image	Changing image takes effect only on VM instances that are newly created or added to the group. The images of existing VM instances do not change.
	Scaling Records	Allows you to view the scaling activities in an auto-scaling group.
	Audit	Audits all of the actions performed on auto-scaling groups , which effectively ensures the security of the Cloud environment.
Snapshot	Create Snapshot	Allows you to schedule snapshot creation at specified time points to record the state of an instance before you perform a business-sensitive operation. This allows rollback in case of breakdowns.
	Manage Snapshot	Displays instances and snapshots on the snapshot management page with instances on the left and snapshots on the right. You can view the relationship between instances and snapshots dynamically.
		The instance panel on the left allows you to sort instances according to their snapshot count or snapshot size.
		The snapshot panel on the right allows you to view the snapshots by list or by topology.
		Allows you to manage the lifecycle of snapshots, such as creating and deleting snapshots.
	Create Instance	Allows you to create an instance from an instance snapshot.
	Revert Snapshot	Allows you to restore an instance from an instance snapshot.
	Audit	Audits all of the actions performed on snapshots, which effectively ensures the security of the Cloud environment.

Type	Features	Description
VM Scheduling Policy	Create VM Scheduling Policy	Allows you to create four types of VM scheduling policies: VM Exclusive from Each Other, VM Affinitive to Each Other, VMs Affinitive to Hosts, and VMs Exclusive from Hosts. The former two define the relationship between VM instances and the latter two define the relationship between hosts and VM instances.
		Every four of the VM scheduling policies can be executed based on either of the following two mechanism: Hard and Soft.
	Manage VM Scheduling Policy	Allows you to manage the lifecycle of VM scheduling policies, such as creating, editing, enabling, disabling, and deleting scheduling policies.
	Associate/Disassociate VM Scheduling Group	Allows you to associate with/disassociate from one or more VM scheduling polices with a VM scheduling group.
	Associate/Disassociate Host Scheduling Group	Allows you to associate/disassociate one or more VM scheduling polices with/from a host scheduling group.
		You can associate/disassociate only VMs Affinitive to Hosts and VMs Exclusive from Hosts with/from a host scheduling group.
	Manage VM Scheduling Group	Allows you to add one or more VM instances to or remove one or more VM instances from a VM scheduling group.
		Allows you to manage the lifecycle of VM scheduling groups, such as creating, editing, and deleting VM scheduling groups.
	Manage Host Scheduling Group	Allows you to add one or more hosts to or remove one or more hosts from a VM scheduling group.
		Allows you to manage the lifecycle of host scheduling groups, such as creating, editing, and deleting host scheduling groups.
	Audit	Audits all of the actions performed on VM scheduling polices, which effectively ensures the security of the Cloud environment.
SSH Key	Create SSH Key	Allows you to generate SSH key pairs on the Cloud or import a generated SSH public key to the Cloud.
		Supported encryption methods: ssh-rsa、ssh-dss、ecdsa-sha2-nistp256、ssh-ed25519、ssh-ecdsa.

Type	Features	Description
	Manage SSH Key	Allows you to manage the lifecycle of SSH keys, such as creating, editing, and deleting SSH keys.
	VM Attachment/ Detachment	Allows you to attach/detach SSH keys to VM instance.
		Allows you to attach one SSH key to one or more VM instance.
Zone	Create Zone	In a data center, a zone corresponds to an equipment room . You can create one or more zones as needed, and create clusters/network resources and primary storage to each zone .
	Manage Zone	Allows you to manage the lifecycle of zones, such as creating , enabling, disabling, and deleting zones.
	Manage Associated Resources	Allows you to manage the clusters, baremetal clusters/elastic baremetal clusters (licenses are required), primary storage, backup storage, L2 networks and other resources in a zone.
	Audit	Audits all of the actions performed on zones, which effectivel y ensures the security of the cloud environment.
Cluster	Create Cluster	Allows you to define cluster attributes (KVM and XDragon) based on the hypervisor type of hosts. Hosts in a KVM cluster use the KVM virtualization technology, and hosts in a XDragon cluster use the X-Dragon architecture.
		Allows you to define cluster attributes based on the host CPU architecture, including x86_64, aarch64, and mips64el.
		Allows you to specify a VDI network and migration network for a cluster.
		Allows you to set the VM CPU model and host CPU model in a cluster as needed.
	Manage Cluster	Allows you to manage the lifecycle of clusters, such as creating, enabling, disabling, and deleting clusters.
	Manage Associated Resources	Allows you to manage the VM instances, hosts, primary storage, iSCSI storage, NVMe storage, L2 networks, peripheral devices, and other resources in a cluster.
	Advanced Settings	Allows you to set the CPU overcommit, memory overcommit, and host reserved memory for all VM instances in a cluster.
		Allows you to enable vNIC multi-queue upgrading for all VM instances in a cluster to improve the VM performance.

Type	Features	Description
		Allows you to enable multi-queue driver support for all VM NICs in a cluster to allocate Virtio NIC traffic to multiple CPUs .
		Allows you to enable huge page for all hosts in a cluster, which effectively reduce the CPU performance loss of VM instances.
		Allows you to enable Hyper-V simulation for all Windows VM instances in a cluster.
		Allows you to set the default graphics card type at the VM startup for all VM instances in a cluster.
		Allows you to enable KVM virtualization flag for all VM instances in a cluster.
		Allows you to enable Dynamic Resource Scheduling (DRS) for clusters. This feature monitors the CPU or memory load of hosts on a cluster basis, and allows you to configure manual or auto DRS strategy to balance cluster loads and improves O&M efficiencies. Manual DRS provides scheduling suggestions based on which you can schedule resources for load balancing. Auto DRS schedules resources based on the system scheduling algorithm without arousing your awareness.
		Allows you to enable Zero Copy for all hosts in a cluster. Enabling this feature will reduce the number of data copies between user space and kernel space, lower CPU usage, and improve vNIC performance.
	Audit	Audits all of the actions performed on clusters, which effectively ensures the security of the cloud environment.
Host	Bulk Action	Allows you to manage hosts in bulk.
	Add Host	Allows you to add hosts manually or by importing a template. You can add up to 500 hosts at a time.
		Supports two hypervisor types: KVM and XDragon. KVM hosts use the KVM virtualization technology and XDragon hosts use the X-Dragon architecture.
	Manage Host	Allows you to manage the lifecycle of hosts, such as adding , enabling, disabling, reconnecting, putting into maintenance mode, deleting, starting, shutting down, and restarting hosts.

Type	Features	Description
	Custom Tag	Allows you to customize tags for hosts so that you can locate them quickly.
	Change Host SSH Password	Allows you to change the SSH password of a host. The new password takes effect after the host automatically reconnects .
	Modify IPMI Info	Allows you to modify the IPMI username and password of a host.
	Enter Web Terminal	Allows you to enter the web terminal of a host and perform operations on the host.
	Add Bond	Allows you to bind multiple physical NICs on the host.
		Supports two bond modes: active-backup mode and LDAP mode.
	Manage Associated Resources	Allows you to manage the VM instances, VPC vRouters, and other virtual resources on a host.
		After you deploy SAN storage (iSCSI storage and FC storage) on a host, you can manage the LUNs on the host and pass through them to VM instances.
		Allows you to manage the physical NICs detected on a host, generate VF NICs from these physical ones through SR-IOV , and pass through the VF NICs to VM instances. These VF NICs inherit the high performance of those physical ones.
		Allows you to manage the physical GPU devices detected on a host and pass through them with other peripheral devices (such as GPU graphics card and GPU sound cards) to VM instances.
		Allows you to generate virtual GPU devices from physical GPU devices (NVIDIA/AMD graphics cards) and attach these virtual GPU devices to VM instances.
		Allows you to manage the USB devices detected on a host and pass through them to VM instances.
		Allows you to manage the PCI devices detected on a host, edit the PCI allowlist, and pass through these PCI devices to VM instances. The PCI devices include Ali-NPU cards, IB cards in PCI mode, and FPGA cards.
		When the overall workload decreases, the Cloud is working on supporting the consolidation of workloads and the

Type	Features	Description
		redistribution of VM instances among hosts in a cluster so that some hosts can be powered off to reduce power consumption.
	Intel EPT Hardware Assist	Allows you to enable Intel EPT hardware assist for Intel CPUs to improve the CPU performance.
	Host Monitoring	Monitors and displays host metrics such as CPU, memory, disk read and write, disk size, and NIC throughput.
	Audit	Audits all of the actions performed on hosts, which effectively ensures the security of the Cloud environment.
	Export CSV File	Allows you to export the host information as a CSV table, which helps in statistical analysis and problem diagnosis.
Primary Storage	Local Storage	Allows you to use the local disk directory of your host as a primary storage.
		Supported backup storage: ImageStore.
		Allows you to manage the lifecycle of local primary storage , such as adding, enabling, disabling, reconnecting, putting into maintenance mode, and deleting local primary storage.
		Allows you to manage VM instances, volumes, clusters, hosts, and other resources on a local primary storage.
		Monitors and displays the percentage of used capacity of the local primary storage.
		Supports predicting physical storage usage trend for local primary storage.
	NFS	Supports NFS protocols. All hosts can automatically mount the same NFS shared directory as the primary storage.
		Supported backup storage: ImageStore.
		Allows you to specify a storage network for NFS primary storage. The storage network is used to check the health status of VM instances.
		Allows you to manage the lifecycle of NFS primary storage , such as adding, enabling, disabling, reconnecting, putting into maintenance mode, and deleting NFS primary storage.
		Allows you to manage VM instances, volumes, clusters, and other resources on a NFS primary storage.

Type	Features	Description
		Allows you to clean up the raw data preserved after migration across NFS primary storage.
		Monitors and displays the percentage of used capacity of NFS primary storage.
		Supports predicting physical storage usage trend for NFS primary storage.
	SharedMountPoint	Supports network shared storage provided by commonly used distributed file systems, such as MooseFS, GlusterFS, OCFS2, and GFS2.
		Supported backup storage: ImageStore.
		Allows you to specify a storage network for SharedMountPoint primary storage. The storage network is used to check the health status of VM instances.
		Allows you to manage the lifecycle of SharedMountPoint primary storage, such as adding, enabling, disabling, reconnecting, putting into maintenance mode, and deleting SharedMountPoint primary storage.
		Allows you to manage VM instances, volumes, clusters, and other resources on a SharedMountPoint primary storage.
		Monitors and displays the percentage of used capacity of SharedMountPoint primary storage.
		Supports predicting physical storage usage trend for SharedMountPoint primary storage.
	Ceph	Supports Ceph distributed block storage. Supported editions: Ceph open source edition (J/L/N) and Ceph enterprise edition.
		If you add Ceph enterprise to the Cloud, you can enjoy the license validity reminder.
		Supported backup storage: ImageStore and Ceph.
		Allows you to specify Ceph pools such as root volume pool, data volume pool, and image cache pool when you add a Ceph primary storage. You can manage all the Ceph pool centrally, add more Ceph pools to expand the capacity, customize the display name of Ceph pool, and specify Ceph pools when you create VM instances, clone VM instances

Type	Features	Description
		, and create volumes. You can also create alarms for Ceph pools.
		Allows you to specify a storage network for Ceph primary storage. The storage network is used to check the health status of VM instances.
		Allows you to add multiple Ceph monitors and manage all the monitors centrally.
		Allows you to manage the lifecycle of Ceph primary storage , such as adding, enabling, disabling, reconnecting, putting into maintenance mode, and deleting Ceph primary storage.
		Allows you to manage VM instances, volumes, block storage volumes, clusters, and other resources on a Ceph primary storage.
		Allows you to clean up the original data preserved after migration across Ceph primary storage.
		Monitors and displays the percentage of used capacity of Ceph primary.
		Supports predicting physical storage usage trend for Ceph primary storage.
	SharedBlock	Allows you to use a block device divided from a SAN storage as a storage pool. SharedBlock primary storage supports iSCSI and FC shared access protocols.
		Supported backup storage: ImageStore.
		Allows you to specify a provisioning method (thick provisioni ng or thin provisioning) when you add a SharedBlock primary storage. You can also specify the provisioning method when you create VM instances, clone VM instances, or create volumes by using a SharedBlock primary storage.
		Allows you to specify a storage network for SharedBlock primary storage. The storage network is used to check the health status of VM instances.
		Allows you to add multiple shared blocks and refresh the storage capacity to view its changes when you expand or replace a block 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 manage the lifecycle of SharedBlock primary storage, such as adding, enabling, disabling, reconnecting, putting into maintenance mode, and deleting SharedBlock primary storage.
		Allows you to manage VM instances, volumes, clusters , LUNs, and other resources on a SharedBlock primary storage.
		Allows you to clean up the original data preserved after migration across SharedBlock primary storage.
		Monitors and displays the percentage of used capacity of SharedBlock primary storage.
		Supports predicting physical storage usage trend for SharedBlock primary storage.
	Support Multiple Primary Storage Per Cluster	Supports more than one local primary storage per cluster.
		Supports more than one NFS primary storage per cluster.
		Supports more than one SharedBlock primary storage per cluster.
		Supports one local primary storage and one NFS/ SharedMountPoint/SharedBlock primary storage per cluster.
		Supports one Ceph primary storage and multiple SharedBlock primary storage per cluster.
		Supports one Ceph primary storage and up to 3 LocalStorage primary storage per cluster.
	Advanced Settings	Allow you to set the space preallocation policy for volumes on local, NFS, SharedMountPoint, and SharedBlock primary storage.
		Allow you to set the storage preallocation policy for SharedBlock primary storage.
		Allow you to set storage overcommit for all types of primary storage.
	Audit	Audits all of the actions performed on primary storage, which effectively ensures the security of the cloud environment.

Type	Features	Description
Backup Storage	ImageStore	Stores image files as image slices and supports incremental storage.
		Supported primary storage: LocalStorage, NFS, SharedMountPoint, Ceph, and SharedBlock.
		Allows you to obtain the existing image files under the mount path of the ImageStore backup storage.
		Allows you to specify a data network for an ImageStore backup storage for data communication with compute nodes.
		Supports image synchronization between different ImageStore backup storage on the same management node, and allows you to specify an image synchronization network for ImageStore backup storage.
		Allows you to manage the lifecycle of ImageStore backup storage, such as adding, enabling, disabling, reconnecting, and deleting ImageStore backup storage.
		Allows you to clean up invalid data stored in ImageStore backup stores to releases storage space.
		Allows you to change the password for an ImageStore backup storage.
		Allows you to centrally manage images in an ImageStore backup storage.
		Monitors and displays the percentage of used capacity of ImageStore primary storage.
	Ceph	Stores image files as Ceph distributed blocks.
		Supported primary storage: Ceph.
		Allows you to add multiple Ceph monitors and manage all the monitors centrally.
		Allows you to specify Ceph pools when you add a Ceph backup storage.
		Allows you to specify a data network for a Ceph backup storage for data communication with compute nodes.
		Allows you to manage the lifecycle of Ceph backup storage , such as adding, enabling, disabling, reconnecting, putting into maintenance mode, and deleting Ceph backup storage.

Type	Features	Description
		Allows you to centrally manage images in a Ceph backup storage.
		Allows you to clean up the original data preserved after migration across Ceph backup storage.
		Allows the deleted data to be retained in the Ceph recycle bin for a period. You can specify this period by setting **Retention Period of Data Deleted from Ceph Primary Storage** in Global Setting. After this period, data is expunged automatically.
		Allows you to manually expunge data in the Ceph recycle bin.
		Monitors and displays the percentage of used capacity of Ceph backup storage.
	Audit	Audits all of the actions performed on backup storage, which effectively ensures the security of the cloud environment.
SAN Storage	iSCSI	Allows you to add an iSCSI server and directly log in to iSCSI storage after you add the server successfully.
		Synchronizes data on iSCSI storage and displays all block devices on iSCSI storage in real time.
		Allows you to add a block device divided from an iSCSI storage as a SharedBlock primary storage and pass through it to VM instances.
		Allows you to manage the lifecycle of iSCSI storage, such as enabling, disabling, and deleting iSCSI storage.
		Allows you to attach/detach an iSCSI storage to/from a cluster.
	FC	Synchronizes device information after you deployed an FC storage and displays the FC storage and its block devices in real time.
		Allows you to add a block device divided from an FC storage as a SharedBlock primary storage and pass through it to VM instances.
		Synchronizes information about a single block device on an FC storage.
		Checks the status of the cluster where block devices are located.

Type	Features	Description
NVMe Storage	/	Synchronizes device information after you deployed an NVMe storage and displays the NVMe storage and its block devices in real time.
		Allows you to add a block device divided from an FC storage as a SharedBlock primary storage.
Physical Network	/	Allows you to attach network-type tags to physical NIC ports to mark the actual usage of the networks they reside on. NIC ports with tags can be displayed on this page by network types or by cluster.
		Allows you to modify the network types of physical NIC ports.
		Allows you to view the flow monitoring based on network types. Three entries are provides: Dashboard, cluster details pages, and host details pages.
Network Resource	L2 Network	Supports the following types of L2 networks: L2NoVlanNetwork, L2VlanNetwork, VxlanNetwork, and HardwareVxlanNetwork.
		VLAN (802 1Q) supports a maximum of 4094 logical networks, and VXLAN supports a maximum of 16 million logical networks.
		VxlanNetwork is a software VXLAN-based solution that effectively addresses the shortage of logical network segments in the cloud data center and MAC flooding in upper layer switches.
		HardwareVxlanNetwork is a solution for working with third-party hardware SDN. By adding an SDN controller, you can take over the SDN network of hardware switches on the Cloud, therefore reducing network latency and improving VXLAN network performance.
		Supports four types of network acceleration mode, including Normal, SR-IOV, and Smart NIC. The normal mode supports all types of L2 networks and the latter two support only L2VlanNetwork and L2NoVlanNetwork.
		Allows you to manage the lifecycle of L2 networks, such as creating and deleting L2 networks.
		Allows you to centrally manage L3 networks and clusters on an L2 network.

Type	Features	Description
	VXLAN Pool	Supports software SDN VXLAN pools and hardware SDN VXLAN pools. A software SDN VXLAN pool is a collection of VxlanNetwork L2 networks, and a hardware SDN VXLAN pool is a collection of HardwareVxlanNetwork L2 networks.
		Allows you to manage the lifecycle of VXLAN pools, such as creating and deleting VXLAN pools.
		Allows you to manage VNI ranges in a VXLAN pool and customize the name of the VNI ranges.
		Allows you to centrally manage the VTEP, clusters, and VXLAN networks in a VXLAN pool.
	Public Network	A public network is an L3 network that has direct access to the Internet.
		Allows you to manage the lifecycle of public networks, such as creating and deleting public networks.
		Allows you to add IP ranges of IPv4 and IPv6 types.
		IPv4 public networks allow you to add either an IP range or an address pool. An address pool can be used to create virtual IP addresses only.
		Allows you to customize the MTU of a public network to limit the size of network transmission packets.
		Monitors and displays the IP usage statistics of public networks, which helps to improve IP planning efficiency.
		Allows you to centrally manage the IP ranges (IPv4/IPv6) and DNS resources on a public network.
	Flat Network	A flat network is an L3 network connected to the network where the host is located and has direct access to the Internet.
		VM instances in flat network networks can use IP resources of an actual network.
		Allows you to manage the lifecycle of flat networks, such as creating and deleting flat networks.
		Allows you to enable or disable IP Address Management for a flat network.
		Allows you to add IP ranges of IPv4 and IPv6 types.

Type	Features	Description
		IPv4 flat networks support the following network services : DHCP, User Data, elastic IP, security group, and port mirroring.
		IPv6 flat networks support the following network services: DHCP, DNS, elastic IP, and security group.
		Allows you to customize the MTU of a flat network to limit the size of network transmission packets.
		Monitors and displays the IP usage statistics of flat networks, which helps to improve IP planning efficiency.
		Allows you to centrally manage the IP ranges (IPv4/IPv6) and DNS resources on a flat network.
	VPC Network	A VPC network is an L3 private network where VM instances can be created. A VM instance in a VPC network can access the Internet through a VPC vRouter.
		Allows you to manage the lifecycle of VPC networks, such as creating and deleting VPC networks.
		Allows you to add IP ranges of IPv4 and IPv6 types.
		IPv4 VPC networks support the following network services : DHCP, User Data, DNS, SNAT, route table, elastic IP , port forwarding, load balancing, IPsec tunnel, security group, dynamic routing, multicast routing, VPC firewall, port mirroring, and netflow.
		IPv6 VPC networks support the following network services: DHCP, DNS, and security group.
		Allows you to attach/detach a VPC vRouter to/from a VPC network.
		Allows you to customize the MTU of a VPC network to limit the size of network transmission packets.
		Monitors and displays the IP usage statistics of VPC networks, which helps to improve IP planning efficiency.
		Allows you to centrally manage the IP ranges (IPv4/IPv6) and DNS resources on a VPC network.
	VPC vRouter	A VPC vRouter is a dedicated VM instance that provides multiple network services.
		Allows you to specify a host on which a VPC vRouter starts.

Type	Features	Description
		Allows you to specify a primary storage when you create a VPC vRouter.
		Allows you to specify a default IPv4 address or IPv6 address for a VPC vRouter.
		Allows you to set a DNS (IPv4/IPv6) on a VPC vRouter and centrally manage all the DNS on the VPC vRouter.
		Allows you to associate the virtual CPUs (vCPUs) of a VPC vRouter with host pCPUs stringently and allocate specific pCPUs for the VPC vRouter, thus improving VPC vRouter performances.
		Allows you to manage the lifecycle of VPC vRouters, such as creating, starting, stopping, restarting, and deleting VPC vRouters.
		Allows you to migrate a VPC vRouter to another host without changing the primary storage. This action is supported only by VPC vRouters in the running state. We recommend that you perform this action during off-peak hours.
		Allows you to migrate a VPC vRouter to another primary storage and host. You can hot migrate a VPC vRouter across different types of primary storage, including LocalStorage ↔SharedBlock, LocalStorage↔NFS, and SharedBlock↔ NFS, or across primary storage of the same type, including SharedBlock↔SharedBlock.
		Allows you to access a VPC vRouter by using a terminal. You can also set the console password for a VPC vRouter.
		Supports auto migration across clusters if you set the parameter **Bind with Cluster** to false. Applicable scenarios: start up a VPC vRouter on another host to achieve HA or migrate a VPC vRouter to another host if the source host enters the maintenance mode.
		Allows you to set the CPU model for a VPC vRouter through three entries: Global Setting, Cluster Setting, and VPC vRouter Setting. The setting takes effect with the following priority: Global Setting < Cluster Setting < VPC vRouter Setting.
		Allows you to enable distributed routing for a VPC vRouter as needed to optimize east-west traffic.

Type	Features	Description
		Allows you to enable the SNAT network service for a VPC vRouter as needed.
		Supports STS to improve network transmission efficiency.
		Supports external monitoring: Collects the VPC vRouter data such as CPU, memory, disk I/O, NIC data from hosts by using libvirt.
		Supports internal monitoring: Collects the VPC vRouter data such as CPU, memory, and disk size data from VM instances by using an agent of the VPC vRouter.
		Allows you to centrally manage the VPC networks, public networks, and DNS resources associated with a VPC vRouter.
		Allows you to set QoS for a VPC vRouter to limit its upstream and downstream bandwidth.
		Allows you to centrally manage the network services provided by a VPC vRouter, such as virtual IP addresses, elastic IP addresses, IPsec tunnels, port forwarding, and load balancing.
		Supports OSPF dynamic routing protocols in large-scale network environment.
		Supports multicast routing to forward multicast messages sent by multicast sources to VM instances.
		Has higher resource priority than VM instances by default. When resource contention occurs, the resource priority is as follows: VM instances with Normal priority < VM instances with High priority < VPC vRouters.
	VPC vRouter HA Group	A VPC vRouter HA group consists of two VPC vRouters. Either VPC vRouter can be a primary or secondary VPC vRouter for the group. If the primary VPC vRouter does not work as expected, the VPC vRouter becomes the secondary VPC vRouter in the group to ensure high availability of business.
		Allows you to manage the lifecycle of VPC vRouter HA groups, such as creating and deleting VPC vRouter HA groups.
		Allows you to add a VPC vRouter to an HA group and centrally manage all VPC vRouters in this group. Any

Type	Features	Description
		configuration changes on a VPC vRouter will apply to its partner VPC vRouter.
	vRouter Image	Supports two types of vRouter images: VPC vRouter image and dedicated-performance LB image.
		Allows you to set the CPU architecture of a vRouter image , including x86_64 and aarch64. Creating VPC vRouters or load balancing instances will inherit the CPU architecture of the original image.
		Allows you to upload a vRouter image by using a URL or local browser.
		Allows you to manage the lifecycle of vRouter images, such as creating, enabling, disabling, deleting, recover, and completely deleting vRouter images.
		Allows you to export a vRouter image on the UI from ImageStore or Ceph backup storage.
		Allows you to centrally manage exported vRouter images and provides the MD5 value of the downloaded image to check the image integrity.
	vRouter Offering	Allows you to manage the lifecycle of vRouter offerings, such as creating, enabling, disabling, and deleting vRouter offerings.
	SDN Controller	Allows you to add external SDN controllers to control network devices such as external switches. This helps to reduce network latency and improve the VXLAN network performanc e.
		Currently, only H3C SDN controllers (VCFC) are supported.
		Allows you to manage the lifecycle of SDN controllers, such as creating and deleting SDN controllers.
	Management Network	A management network is used to manage physical resources in the Cloud.
		Allows you to manage the lifecycle of management networks , such as creating and deleting management networks.
		Allows you to add IP ranges of the IPv4 type.
		Allows you to customize the MTU of a management network to limit the size of network transmission packets.

Type	Features	Description
		Monitors and displays the IP usage statistics of management networks, which helps to improve IP planning efficiency.
		Allows you to centrally manage the IP ranges (IPv4) on a management network.
	Flow Network	A flow network is a dedicated network for port mirror transmission. You can use a flow network to transmit the mirrors of data packets of NIC ports to the target ports.
		Allows you to manage the lifecycle of flow networks, such as creating and deleting flow networks.
		Allows you to add IP ranges of the IPv4 type.
		Monitors and displays the IP usage statistics of flow networks , which helps to improve IP planning efficiency.
		Allows you to centrally manage the IP ranges (IPv4) on a flow network.
	Audit	Audits all of the actions performed on network resources , which effectively ensures the security of the cloud environment.
Network Service	Security Group	Provides security control over VM instances on L3 networks.
		Allows you to manage the lifecycle of security groups, such as creating, enabling, disabling, and deleting security groups.
		Allows you to add/delete ingress/egress rules to/from a security group and manage these rules centrally, including modifying,enabling, disabling, importing, and exporting rules and setting rule priorities..
		Supports ALL, TCP, UDP, and ICMP protocols for security group rules.
		Allows you to set a source security group by security group rules.
		Security group rules apply the allowlist and blocklist mechanism.
		For newly created security groups, ingress and egress rules with the ALL protocol type are configured by default. The rules allow mutual communications among VM instances in the same security group.

Type	Features	Description
		Allows you to centrally manage VM NICs associated with a security group.
	Virtual IP Address (VIP)	Provides multiple network services by using VIPs in bridged network environments.
		Divides VIPs into public VIP, flat network VIP, and VPC VIP based on the network where the VIP is created.
		Divides VIPs into system VIP and custom VIP based on how the VIP is created.
		Allows you to manage the lifecycle of VIPs, such as adding and deleting VIPs.
		Allows you to set QoS for public VIPs and flat network VIPs.
		Monitors and displays VIP metrics such as network traffic and network packet rate.
	Elastic IP Address (EIP)	IP addresses in a private network are translated into an EIP that is in another network. This way, private networks can be accessed from other networks by using EIPs.
		Divides EIPs into public EIP and flat network EIP based on the network where the EIP is created.
		Allows you to manage the lifecycle of EIPs, such as adding and deleting EIPs.
		Allows you to associate/disassociate an EIP with/from a VM NIC.
		Allows you to change the owner of an EIP.
	Port Forwarding	Works based on the layer-3 forwarding service provided by VPC vRouters and forwards traffic flows of specified IP addresses and ports in a public network to the specified ports of VM instances. If your public IP addresses are insufficient, you can configure port forwarding for multiple VM instances by using one public IP address and port.
		Supports TCP and UDP.
		Allows you to manage the lifecycle of port forwarding, such as creating and deleting port forwarding.
		Allows you to associate/disassociate port forwarding with/ from a VM NIC.

Type	Features	Description
	Load Balancing	Distributes traffic flows of a VIP to backend servers. It automatically inspects the availability of backend servers and isolates unavailable servers during traffic distribution, which improves the availability and service capability of your business.
		Supports two types of load balancing services: shared-performance load balancing that works based on VPC vRouters and dedicated-performance load balancing that works based on load balancer instances.
		Allows you to manage the lifecycle of load balancers, such as creating and deleting load balancers.
		Allows you to create shared-performance load balancers by using public networks or VPC networks.
		Allows you to create dedicated-performance load balancers by using public networks, flat networks, or VPC networks.
		Monitors and displays load balancer metrics such as inbound /outbound traffic and active/concurrent/new sessions.
		Allows you to centrally manage listeners, backend server groups, and other resources associated with load balancers.
		Allows you to manage the lifecycle of load balancers, such as creating and deleting load balancers.
		Listener protocols support TCP, HTTP, HTTPS, and UDP.
		Supports multiple load balancing algorithms, including Round Robin, Min Connections, Source IP Hash, and Weighted Round Robin.
		Health check protocols support TCP, HTTP, and UDP.
		Listeners that use the HTTPS protocol allow you to associate /disassociate certificates. You can upload certificates or certificate chains and manage these certificates centrally.
		Listeners that use the HTTP or HTTPS protocol allow you to configure forwarding rules for domain forwarding and manage these rules centrally.
		Supports two session persistence mechanisms: TCP/UDP -based 4th-layer session persistence and HTTP/HTTPS-based 7th-layer session persistence

Type	Features	Description
		4th-layer session persistence uses Source IP Hash algorithm to direct requests from clients of the same source IP address to a backend server.
		7th-layer session persistence supports Round Robin, Weighted Round Robin, and Weighted Round Robin. When using Round Robin or Weighted Round Robin algorithm, a load balancer inserts or rewrites a cookie to direct requests to the backend server previously responded. When using Source IP Hash algorithm, a load balancer uses the Hash function to direct requests from clients of the same source IP address to a backend server.
		Allows you to associate/disassociate listeners with/from backend server groups.
		Listeners that use a weighted round-robin load-balancing algorithm allow you to set the weight value for each backend server separately.
		Allows you to manage the lifecycle of backend server groups , such as creating and deleting backend server groups.
		Allows you to add/remove backend server to/from backend server groups.
		Allows you to add VM NICs or other servers outside of the Cloud as backend servers. Note that the later method applies to only dedicated-performance load balancers.
		Dedicated-performance load balancers allows you to create/ delete load balancer offerings.
	VPC Firewall	Monitors ingress and egress traffic of VPC vRouters and decides whether to allow or block specific traffic based on a defined set of security rules.
		Allows you to manage the lifecycle of VPC firewalls, such as creating and deleting VPC firewalls.
		Allows you to centrally manage rules and rule sets associated with VPC firewalls.
		Allows you to manage the lifecycle of rules, such as adding, enabling, disabling, and deleting rules.

Type	Features	Description
		Configures ingress and egress rules by default after a VPC firewall is created and allows you to customize these rules as needed.
		Allows you to manually add rules to a VPC firewall by specifying a single IP address or an IP/port set.
		Allows you to add multiple rules to a VPC firewall by importing a template. You can also modify the rule template and upload it as needed.
		Allows you to set priorities for VPC firewall rules.
		VPC firewall rules have three behaviors: Accept, Drop, and Reject.
		VPC firewall rules support the following packet status: new (new connection requests), established (established connections), invalid (unidentifiable connections), and related (new connection requests that are associated with existing connections).
		VPC firewall rules support the following protocols: ALL, TCP, UDP, ICMP, GRE, ESP, AH, IPIP, VRRP, IPENCAP, PIM, OSPF, and IGMP.
		Allows you to manage the lifecycle of rule sets, such as adding and deleting rule sets.
		Allows you to centrally manage rules and network resources in a rule set.
		Modifications on rules in a rule set take effect after you synchronize the modifications.
		Allows you to save firewall rules as a rule template (managed by the Cloud or export them as a CSV file (offline).
		Allows you to manage the lifecycle of rule templates, such as creating and deleting rule templates.
		Allows you to save IP/port sets as a generic template.
		Allows you to manage the lifecycle of IP/port sets, such as adding, enabling, disabling, and deleting IP/port sets.
	IPsec Tunnel	Encrypts and verifies IP packets that transmit over a virtual private network (VPN) from one site to another.

Type	Features	Description
		IPsec negotiation mode: Supports only the Main mode due to security reasons. The Aggressive mode is not supported.
		IPsec IKE configurations: Support IKEv1 and IKEv2(default)
		IPsec security protocol: Supports only the Encapsulating Security Payload (ESP) protocol.
		IPsec encapsulation mode: Supports only the Tunnel mode. The Transport mode is not supported.
		IPsec routing model: Supports only policy-based IPSec VPN . Route-based IPSec VPN is not supported. Therefore, the tunnel supports only unicast data, and does not support multicast and broadcast.
		Allows you to manage the lifecycle of IPsec tunnels, such as creating and deleting IPsec tunnels.
		Monitors the connection status of IPsec tunnels
		Allows you to centrally manage network resources associated with an IPsec tunnel.
	Dynamic Routing	Supports Open Shortest Path First (OSPF) protocols.
		Allows you to manage the lifecycle of OSPF areas, such as creating and deleting OSPF areas.
		Supports two types of OSPF areas: Standard and Stub.
		Provides three authentication methods for OSPF areas: None , Plaintext, and MD5.
		Allows you to centrally manage the routing configuration of OSPF areas.
	Netflow	Monitors the ingress and egress traffic of the NICs of VPC vRouters.
		Allows you to manage the lifecycle of netflows, such as creating and deleting netflows.
		Supports two versions of data flows: V5 and V9.
		Allows you to centrally manage the routing configuration of netflows.
	Port Mirroring	Mirrors the traffic data of VM NICs and sends the traffic data to the target ports. This helps to analyze the data packets of

Type	Features	Description
		ports, which simplifies the data monitoring and management and makes it easier to locate network errors and exceptions.
		Allows you to manage the lifecycle of port mirroring, such as creating, enabling, disabling, and deleting port mirroring.
		Supports three types of sessions: Ingress, Egress, and Bi-direction.
		Allows you to centrally manage port mirroring sessions.
	Route Table	Allows you to customize routing configurations as needed.
		Allows you to manage the lifecycle of route tables, such as adding, enabling, disabling, and deleting route tables.
		Allows you to centrally manage route entries and VPC vRouter resources in routing tables.
		Allows you to manage the lifecycle of route entries, such as adding and deleting route entries.
		Supports two types of route entries: static route entries and blackhole route entries.
		Allows you to set route priorities.
	Multicast Route	Forwards the multicast messages sent by the multicast source to VM instances, which realizes point-to-multipoint connection between the sender and the receiver.
		Allows you to enable multicast route as needed. After enabled, the multicast route takes effect for all networks associated with VPC vRouters.
		Supports PIM-SM and PIM-SSM routing protocols. In the PIM -SM protocol, RP routers are the essential device in the PIM -SM domain. The RP addresses support static configuration and dynamic election through the BSR mechanism.
		Allows you to centrally manage the multicast configuration tables and multicast routing tables.
	Audit	Audits all of the actions performed on network resources , which effectively ensures the security of the cloud environment.
CloudFormation	Resource Stack	Allows you to manage the lifecycle of resource stacks, such as creating and deleting resource stacks.

Type	Features	Description
		Allows you to create resource stacks by using a stack template (system template or custom template), uploading a file (in UTF8-encoded format), or customizing a text (in the designer).
		Allows you to preview the template configurations before you complete the creation.
		Allows you to centrally manage the templates, resources, and events associated with a resource stack.
	Stack Template	Allows you to manage the lifecycle of stack templates, such as creating, enabling, disabling, and deleting stack templates.
		Allows you to create stack templates by customizing a text or uploading a file.
		Allows you to modify the template content as needed.
	Sample Template	Provides commonly used sample templates for your reference.
		Allows you to manage the lifecycle of sample templates, such as enabling and disabling sample templates.
	Designer	Allows you to establish dependencies between resources by drag-and-drop connections on the canvas.
		Allows you to undo, redo, zoom in, zoom out, fit to canvas, delete, and clear the canvas.
		Allows you to set global parameters of the following types: String, Number (integer or floating point), Comma-delimited list (equivalent to List<String> in Java), and Boolean.
		Allows you to preview templates, generate resource stacks, and save as stack templates.
	Audit	Audits all of the actions performed on CloudFormation, which effectively ensures the security of the cloud environment.
Network Topology	Global Topology	Displays the network topology in the Cloud, helping you to manage and maintain your networks more efficiently.
		Allows you to refresh the topology to view latest information.
		Allows you to export the global topology in PNG format.
		Allows you to hide or unhide VM instances, highlight the selected resources, view the resource information in hover, and view the VM/VPC vRouter status.

Type	Features	Description
		Allows you to fit to window and zoom in, zoom our the canvas .
		Allows you to search for resources by resource category and attribute within the current global topology.
	Custom Topology	Allows you to generate a custom topology.
		Allows you to refresh the topology to view latest information.
		Allows you to export the custom topology in PNG format.
		Allows you to highlight the selected resources, view the resource information in hover, and view the VM/VPC vRouter status.
		Allows you to fit to window and zoom in, zoom our the canvas .
		Allows you to search for resources by resource category and attribute within the current global topology.
Performance Analysis	View Performance Analysis	Displays the performance metrics of key resources.
		Allows you to view the monitoring data by resources, including VM instance, VPC vRouter, host, backup storage, L3 network, and virtual IP.
		Supports two monitoring methods: external monitoring and internal monitoring.
		Allows you to view the monitoring data by selecting a time span. Available time spans: 15 minutes, 1 hour, 1 week, and custom.
		Supports advanced filtering, including filter by monitoring items (metrics and thresholds), resource scope (all resources /specified resources), and owner scope (all owners/specified owners).
		Allows you to sort the items by resource name or monitoring metric.
		Allows you to view the monitoring data details of a single resource.
		Allows you to customize the number of items to be displayed on each page. By default, 10 items are displayed per page.
	Export Performance Analysis Report	Allows you to export all the report information or export the information on the current page in CSV format.

Type	Features	Description
		Allows you to export the average, maximum, or minimum values of the metrics for VM instances and VPC vRouters.
Capacity Management	Resource Capacity Card	Displays the capacities and usages of key resources as cards in the Cloud.
		Supports the following resources: primary storage, backup storage, management node, VM instance, volume, image, snapshot, and compute node.
		Allows you to jump to the corresponding resource list from the current card.
	Resource Capacity Top 10	Allows you to view top 10 resources based on the capacity usage.
		Supports the following resources: host, primary storage, backup storage, VM instance, volume, image, and snapshot.
		Allows you to sort resources by capacity utilization, used physical capacity, available physical capacity, total physical capacity.
		Allows you to view the disk usage details of a single resource .
Management Mode Monitoring	/	Allows you to view the health status of each management node in a multi-management node environment.
		Allows you to view the management IP and node status.
		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.
Monitoring and Alarm	Alarm	Monitors time-series data and events and sends alarm messages to specified endpoints.
		Supports default alarms and custom alarms.
		Supports resource alarms, event alarms, and extended alarms.
		Allows you to manage the lifecycle of default resource alarms , such as enabling and disabling default resource alarms.
		Allows you to manage the lifecycle of custom resource alarms, such as enabling and disabling custom resource alarms.

Type	Features	Description
		Allows you to create resource alarms for two types of time-series data: resource utilization and resource capacity.
		Provides three emergency levels for resource alarms: 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.
		Allows you to centrally manage the endpoints and alarm records of a resource alarm.
		Allows you to manage the lifecycle of custom event alarms , such as creating, deleting, enabling and disabling custom event alarms.
		Provides three emergency levels for event alarms: emergent , major, and info.
		Allows you to centrally manage the endpoints and alarm records of an event alarm.
		Allows you to manage the lifecycle of extended alarms, such as enabling and disabling extended alarms.
		Allows you to centrally manage the endpoints and alarm records of an extended alarm.
	One-click Alarm	Provides a set of alarm rules for critical resources, which can be used to quickly establish monitoring and alarm services for these resources.
		Applies to resources such as hosts, VM instances, and VPC vRouters.
		Allows you to enable or disable one-click alarms.
		Allows you to enable, disable, and modify a single alarm rule for a one-click alarm.
	Alarm Template	Encapsulates alarm rules as a template and works with resource groups. You can configure alarm rules for resources in bulk, which helps to improve the O&M efficiency.
		Allows you to manage the lifecycle of alarm templates, such as creating and deleting alarm templates.

Type	Features	Description
		Allows you to add/remove alarm rules to/from alarm templates and centrally manage these rules in an alarm template.
		Allows you to attach/detach tags to/from an alarm template.
		Allows you to clone an alarm template.
		Allows you to associate/disassociate resource groups with/ from an alarm template and centrally manage these resource groups of an alarm template.
	Resource Group	Groups resources based on business requirements and works with alarm templates. You can configure alarm rules for resources in bulk, which helps to improve the O&M efficiency.
		Allows you to manage the lifecycle of resource groups, such as creating and deleting resource groups.
		Allows you to add/remove resources to/from a resource group and centrally manage these resources in a resource group.
		Allows you to attach/detach tags to/from a resource group.
		Allows you to associate/disassociate alarm templates with/ from a resource group.
		Allows you to centrally manage the alarms, endpoints, and alarm records of a resource group.
	Message Template	Sends messages to endpoints by using a text template.
		Allows you to manage the lifecycle of message templates, such as creating and deleting message templates.
		Supports the following endpoints: email, DingTalk, Microsoft Teams, and short message.
		Supports the following alarm types: resource alarm and event alarm.
		Supports the following types of message texts: alarm message text and recovery message text.
		Allows you to make a template default or cancel the default setting. Only one default template is allowed.
		Allows you to modify the content in a message template.

Type	Features	Description
	Message Source	Allows you to connect to extended message sources.
		Allows you to manage the lifecycle of message sources, such as creating, enabling, disabling, and deleting message sources.
		Supports Ceph Enterprise.
		Provides preconfigured alarm message conversion template and allows you to customize parameters in the template.
	Endpoint	Allows you to obtain your subscribed information by using an endpoint.
		Supports default endpoints and custom endpoints (email, short message, HTTP application, DingTalk, Microsoft Teams , and SNMP trap receiver).
		Default endpoints receive messages sent from the Cloud.
		Allows you to manage the lifecycle of default endpoints, such as enabling and disabling default endpoints.
		Allows you to manage the lifecycle of custom endpoints, such as creating, enabling, disabling, and deleting custom endpoints.
		Allows you to add/remove alarms to/from an endpoint and centrally manage these alarms, including resource alarms and event alarms.
		Allows you to centrally manage messages (alarm messages and extended messages) received by an endpoint.
	Audit	Audits all of the monitoring and alarm actions, which effectively ensures the security of the cloud environment.
Alarm Message	Cloud Platform Alarm Message	Allows you to view and centrally manage alarm messages sent from the Cloud .
		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.
		Allows you to filter messages by resource.

Type	Features	Description
		Allows you to filter messages by specifying a time span.
		Allows you to mark alarm messages as read and filter read or unread messages as needed.
		Allows you to filter messages by emergency levels (emergent, major, and info).
		Allows you to filter messages by alarm type (resource alarm and event alarms).
		Allows you to converge and sort alarm messages based on the alarm times.
		Allows you to set a silence period for alarm messages. During the silence period, no alarm 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.
	Extended Alarm Message	Allows you to take over third-party alarm messages and push, manage them centrally.
		Allows you to mark alarm messages as read and filter read or unread messages as needed.
		Allows you to filter messages by specifying a time span.
One-click Inspection	Five Inspection Categories	Provides five inspection categories, including platform, compute, network, storage, and global setting. These categories cover all key resources and services of the Cloud.
	Multi-layer Healthiness Scoring Mechanism	Provides an in-built three-layer healthiness scoring mechanism that scores resources and services, inspection items, and the overall Cloud. It also displays the score of healthiness for the overall Cloud.
	O&M Suggestion	Provides O&M suggestions on resources in warning or fault status.
	Inspection Report	Provides inspection introduction, summary, and results, and details of abnormal inspection items as well as O&M suggestions.

Type	Features	Description
	Inspection Management	Allows you to select inspection items for one-click inspection.
		Allows you to pause, resume, and cancel inspection, implement re-inspection, and export PDF-formatted inspection reports.
Operation Log	Current Task	Allows you to view and manage operations that are being performed.
		Displays the task progress and remaining time in real time.
		Allows you to cancel, suspend, and continue a current task as needed.
		Allows you to view the details about a current task.
	Historic Operation	Displays the historic operations performed in the Cloud.
		Allows you to view all the operations that were performed.
		Allows you to filter operation logs by specifying a time span.
		Allows you to filter operation logs by task results, including succeeded, failed, canceled, canceling, exception, timeout, suspended, and unknown.
		Allows you to filter operation logs by operators.
		Allows you to export operation logs in CSV format.
		Allows you to view the details about an operation log.
		Allows you to set the operation log retention period in the Global Setting.
	Auto-Scheduling Logs	Displays the VM auto-scheduling logs triggered by the management node, such as VM recovery from HA and host maintenance.
		Allows you to view all the auto-scheduling logs that were triggered.
		Allows you to filter auto-scheduling logs by specifying a time span.
		Allows you to filter auto-scheduling logs by task results, including succeeded and failed.
		Allows you to export auto-scheduling logs in CSV format.
		Allows you to view the details about an auto-scheduling log.

Type	Features	Description
		Allows you to set the auto-scheduling log retention period in the Global Setting.
Audit	/	Monitors and records all activities in the Cloud, which effectively ensures the security of the cloud environment.
		Allows you to filter audit records by resource actions and login actions.
		Allows you to filter audit records by specifying a time span.
		Allows you to filter audit records by task results, including succeeded and failed.
		Allows you export audit records in CSV format.
		Allows you to view the details about an audit record.
Log Collection	Collect Log	Allows you to collect the logs of the Cloud and of various nodes on the Cloud that are generated in the specified time range.
	Manage Log	Allows you to collect, recollect, download, delete, and cancel the collection of logs.
Scheduled O&M	Scheduled Job	Allows you to manage the lifecycle of scheduled jobs, such as creating, enabling, disabling, and deleting scheduled jobs.
		Supports VM instances and volumes.
		Allows you to view job records centrally.
		Allows you to attach/detach schedulers to/from a scheduled job.
	Scheduler	Allows you to manage the lifecycle of schedulers, such as creating and deleting schedulers.
		Allows you to centrally manage the scheduled jobs of a scheduler.
		Allows you to centrally manage schedulers that were completed.
	Audit	Audits all of the scheduled O&M actions, which effectively ensures the security of the cloud environment.
Tag	/	Allows you to customize tags for resources and quickly locate resources by tag type and tag name.
		Supports admin tags and tenant tags.

Type	Features	Description
		Allows you to manage the lifecycle of tags, such as creating and deleting tags.
		Allows admins to attach/detach tags to/from all resources on the Cloud and tenants to attach/detach tags to/from resources of tenants.
		Allows you to centrally manage resources with a tag attached .
		Audits tag actions, which effectively ensures the security of the cloud environment.
Billing Management	Bills	A bill is the expense of resources totaled at a specified time period. Billing is accurate to the second.
		Supported bill types; project bills, department bills, and sub-account bills.
		Allows you to filter bills by specifying a time span.
		Allows you to view project bills in a list, export all of the project bills in CSV format, view the billing details of a single project, and export the bills of a single project in CSV format.
		Allows you to view department bills in a list, view the bills of the current department or its sub-departments, view the bills of directly affiliated projects in a list, export total bills of all directly affiliated projects in CSV format, view the billing details of a single project, and export the bills of a single project in CSV format.
		Allows you to view sub-account bills in a list, export all of the sub-account bills in CSV format, view the billing details of a single sub-account, and export the bills of a single sub-account in CSV format.
		Allows you to disable the billing feature in Global Setting. Then, the system stops billing resources and bills are no longer generated.
		By default, bills are generated every day at 00:00. You can change the bill generation time in the Global Setting.
		Allows you to set the currency symbol displayed on the UI in the Global Setting. Default value: ¥. Valid values: ¥, $, €, £, A $, HK$, ¥, CHF, and C$.

Type	Features	Description
	Pricing List	A pricing list is a list of unit prices of different resources. The unit price of a resource is set based on the specification and usage time of the resource.
		Allows you to manage the lifecycle of pricing lists, such as creating and deleting pricing lists.
		Allows you to set the unit price for the following resources : CPU/memory, volume (root volume/data volume), GPU device (desktop GPU and compute GPU), network (VM public IP and virtual IP), and elastic baremetal instance (elastic baremetal offering).
		Allows you to generate bills based on disk performances. You can set the billing unit price for root volumes and data volumes with different performances by setting advanced parameters.
		Allows you to modify the billing unit price as needed.
		Allows you to centrally manage the price history and related resources.
Access Control	Console Proxy	Allows you to set a console proxy to log in to a VM instance.
		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 Cloud.
		Supports two types of AccessKey: local AccessKey and third -party AccessKey.
		Allows you to manage the lifecycle of local AccessKeys, such as generating, enabling, disabling, and deleting local AccessKeys.
		Allows you to manage the lifecycle of third-party AccessKeys , such as generating and deleting third-party AccessKeys.
		Audits all of the AccessKey actions, which effectively ensures the security of the cloud environment.
	IP Allowlist/Blocklist	An IP blocklist or allowlist identifies and filters IP addresses that access the Cloud.
		Allows you to enable the IP allowlist/blocklist feature in the Global Setting as needed.

Type	Features	Description
		Allows you to manage the lifecycle of IP allowlists/blocklists, such as adding and deleting IP allowlists/blocklists.
		Audits all of the IP allowlist/blocklist actions, which effectively ensures the security of the cloud environment.
Application Center	/	Allows you to add URLs of third-party applications. This allows you to manage the applications in a centralized way and quickly open the applications.
		Supports the following types of applications: storage, database, security, IaaS, PaaS, and SaaS applications.
		Allows you to set the sharing mode of a resource, including share globally, share to specified projects or accounts, and not share.
		Allows you to manage the lifecycle of applications, such as adding and deleting applications.
Sub-Account Management	/	A sub-account is created by the admin or synced from a third-party authentication system and is managed by the admin. Resources created under a sub-account are managed by the sub-account.
		Allows you to manage the lifecycle of local sub-accounts, such as creating and deleting local sub-accounts.
		Allows you to add a 3rd-party authentication server to the Cloud so as to integrate the 3rd-party authentication system and enable password-free login of related accounts in the system.
		The supported authentication server type includes OIDC.
		Allows you to configure user mapping rules for the OIDC server.
		Allows you to manage the lifecycle of the 3rd-party authentication server, such as adding and deleting the 3rd-party authentication server.
		Allows you to manage the lifecycle of 3rd-party sub-accounts, such as synchronizing and deleting 3rd-party sub-accounts.
		Allows you to set the initial password or change the password of a sub-account.

Type	Features	Description
		Allows you to bill for resources used by sub-accounts, attach pricing lists to a sub-account, and change pricing lists for the sub-account.
		Allows you to set two-factor authentication for sub-account login, view the two-factor QR codes of the sub-account, and download the two-factor QR codes.
		Allows you to set and manage resource quota for sub-accounts, including compute resources, storage resources, and network resources.
		Allows you to centrally manage the associated or shared resources of a sub-account.
		Audits all of the sub-account actions, which effectively ensures the security of the cloud environment.
System Setting	Theme and Appearance	Allows you to customize the theme and appearance of the Cloud.
		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.
	Time Management	Allows you to configure NTP time servers for the Cloud to sync the clock of the time servers with all nodes of the Cloud . Three time protocol modes are supported: Internal, Internal and External, and External.
		Allows you manually sync time by force to save your time.
		Displays the latest system UTC date, time, and time zone.
	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.
		Allows you to manage the lifecycle of email servers, such as adding, enabling, disabling, and deleting email servers.
		Supported email server type: SMTP.
		Supported encryption type: STARTTLS, SSL/TLS, and NONE.
		Allows you to test the email server connectivity.
		Allows you to change the owner of email servers.

Type	Features	Description
		Audits all of the email server actions, which effectively ensures the security of the cloud environment.
	Log Server	A log server is used to collect logs of the management node . You can add a log server to the cloud and use the collected logs to locate errors and exceptions. This improves your O& M efficiency.
		Allows you to manage the lifecycle of log servers, such as adding and deleting log servers.
		Allows you to set the log severity from LOCAL0 to LOCAL7.
		Allows you to test the log server connectivity.
		Audits all of the log server actions, which effectively ensures the security of the cloud environment.
	SNMP Management	Connects 3rd-party platform and Cloud through SNMP, enabling the 3rd-part platform to get monitoring data from Cloud or receive alarms pushed from Cloud.
		Allows you to enable/disable SNMP Management.
		Allows you to configure SNMP parameters in a visual method .
		Allows you to add SNMP trap receivers to receive alarms from Cloud.
		Allows you to add SNMP trap receivers as endpoints and attach them to specified alarms.
Platform Setting	HA Policy	HA Policy is a mechanism that ensures sustained and stable running of the business if VM instances are unexpectedly or scheduled stopped or are errored because of errors occurred to compute, network, or storage resources associated with the VM instances.
		Provides None and NeverStop VM HA modes, which specify whether to enable auto restart if VM instances are stopped.
		Allows you to configure VM Failover Strategy in a table based on the management network connectivity status, storage network connectivity status, and business NIC status.
		Allows you to modify host error detection settings and advanced HA-related settings. These settings take effect on the Cloud.

Type	Features	Description
		Allows you to view and filter VM HA logs.
	Scenario Template	Provides multiple templates that encapsulate scenario-based global settings. You can apply a template globally with one click based on your business needs. This improves your O&M efficiency.
		Applies to VM performance optimization, restoration from high availability, cloud security setting, and production environment setting.
		Allows you to apply a scenario template with one click.
		Allows you to reset to default settings with one click.
		Allows you to modify settings of a single item in a scenario template.
	Global Setting	Allows you to configure settings that take effect on the whole platform.
		Support basic settings and advanced settings.
		Allows you to reset to default settings with one click.
		Supports quick search and directory navigation to help you quickly locate target items.
		Allows you to modify settings of a single item in the Global Setting.
System and Security	Version Detection	Allows you to enable version detection which periodically detects the latest version including production environment recommended versions or technical preview versions
		Allows you to specify the auto detection duration by day, week, month, or year.
		Allows you to implement manual detection or use the auto detection. It provides information about the version number and the highlights if the latest version is available.
	Experience Improvement Program	Allows you to join in or opt out the Experience Improvement Program.
	Certificate Management	Allows you to configure and manage a SSL certificate, including third-party certificate and system self-signed certificate.
License Management	/	Licensing in the Cloud is supplied in different functionality packages as Base and Plus.

Type	Features	Description
		You can purchase a package as needed.
		The Base license provides the basic and essential features of the Cloud, which can meet the mainstream business requirements.
		Functionalities covered in the Base license include Standard, Enterprise Trial, and Enterprise Prepaid.
		The Plus license provides add-on features or feature enhancements to meet the specific business requirements.
		Functionalities covered in the Plus license include VMware Management, Tenant Management, ARM64 Management , Backup Service, Continuous Data Protection (CDP) Service, Migration Service, Baremetal Management, Elastic Baremetal Management, Alibaba Cloud Hybrid Cloud Management, Cryptography Security Compliance, 5x8 (7x24) After-Sales Service, SR-IOV NIC Service, GPU Service, Billing Management, CloudFormation, Auto-Scaling Service, and Smart NIC Service.
		A Base license is required to install the Plus license.
		Supports two licensing methods: USB key and request key.
		The USB key licensing method allows you to obtain the authorization by inserting only one USB key into the management node.
		The request key method allows you to obtain the authorizat ion by uploading the license file to the management node.
		Allows you to view the current license status and licensing records.
		Allows you to delete a Plus license as needed.
		Provides license expiration reminders when your license is about to expire, expired, or license quota exceeds.
Cloud Login	Login Method	Allows you to access the UI via HTTP or HTTPS.
		Supports account login and tenant login.
		Allows you to access the Cloud and experience all of the features by using command lines.

Type	Features	Description
	Login Security	Allows you to set the maximum number of continuous login failures that trigger verification by verification code. Default: 6 .
		Supports two-factor authentication, which further enhances the account security.
		Allows you to set the login password complexity by set the password length and characters combined of digits, uppercase/lowercase letters, and special characters.
		Allows you to set the password validity period by customizing the password update cycle. We recommend that you change the login password regularly to ensure the login security.
		Supports historical password check and allows you to customize the number of recent passwords that cannot be reused.
		Allows you to specify whether to lock the login account if the logins continuously fail, the number of allowed failed attempts , and how long the account will be locked.
		Allows you to specify whether to disallow simultaneous connection sessions established by one user. If yes, one user can establish only one connection session with the platform.
		Allows you to set the login interface with the default link.
VDI	Solution	Supports SPICE, RDP, and VNC protocols.
		Allows you to specify a VDI network.
		Supports USB redirection, which means multiple USB devices are compatible.
		Allows you to set an independent VDI network.
		Supports multi-screen display.
		Supports microphones.
		Supports SPICE to optimize traffics.
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, resources, and documents.

Type	Features	Description
	Embedded Document	Provides embedded documents in the help center.
Installation	/	Allows you to complete installing and deploying the Cloud from scratch within just 30 minutes with one simple command .
		Supports the following installation modes: Tenant Management Mode, Community Management Mode, Compute Node Mode, Expert Mode, and Simplified Expert Mode.
		Supports ISO: h76c ISO, h79c and h84r ISO.
		Allows you to burn ISO images to U drives by using Rufus.
Upgrade	Seamless Upgrade	Allows you to seamlessly upgrade the Cloud from an earlier version to a later version.
	Deployment Environment Upgrade	Allows you to specify the deployment environment from the Expert Mode.

Features in VMware Management

Type	Features	Description
vCenter	Basic Resource	Allows you to take over vCenter 5.5, 6.0, 6.5, 6.7, and 7.0.
		Supported protocols: HTTPS (default) and HTTP.
		Supports automatic and manual data synchronization. Automatic data synchronization occurs when a vCenter is added to the Cloud for the first time. You can also enable **vCenter Data Auto Sync** in the Global Setting and set an automatic synchronization interval to realize a regular automatic data synchronization.
		Allows you to centrally manage resources associated with a vCenter, including clusters, primary storage, backup storage , hosts, and resource pools.
		Allows you to delete a taken over vCenter from the Cloud . This deletes only the local record of the vCenter and associated resources but does not affect the real resources in the remote vCenter.
	VM Instance	Allows you to manage the lifecycle of vCenter VM instances , such as creating, booting, stopping, rebooting, resuming, pausing, powering off, deleting, and recovering vCenter VM instances.

Type	Features	Description
		Allows you to launch the console of a vCenter VM instance and set the console password as needed.
		Allows you to clone a vCenter VM instance online or offline without data volumes.
		Allows you to hot migrate a vCenter VM instance across shared primary storage with data volumes attached.
		Allows you to modify the instance offering (CPU/memory) of a stopped vCenter VM instance.
		Allows you to change the owner of a running or stopped vCenter VM instance.
		Allows you to set the HA level (None/NeverStop) for a vCenter VM instance. You can enable **VM HA** in the Global Setting as needed.
		Allows you to attach custom tags to vCenter VM instances for an efficient resource location.
		Supports an external monitoring on the CPU, memory, disk, virtual disk, and NIC of a vCenter VM instance.
		Allows you to centrally manage resources associated with a vCenter VM instance, such as volumes and NICs.
	Network	Supported L2 networks: L2NoVlanNetworks and L2VlanNetworks.
		Supported L3 networks: public networks, flat networks, and VPC networks.
		Supported switch types: dvSwitch and vSwitch.
		Supported VPC network services: SNAT, DHCP, elastic IP, port forwarding, load balancing, and IPsec tunnel.
		Allows you to manage the lifecycle of networks, such as creating L2/L3 networks and deleting L3 networks.
		Allows you to set the sharing mode for an L3 network. Valid values: share globally, share to specified projects or accounts, and not share.
		Provides a list displaying the IP usage of an L3 network to improve IP planning efficiency.
		Allows you to attach a cluster to the L2 network an L3 network belongs to.

Type	Features	Description
		Allows you to centrally manage IPv4 network ranges of an L3 network.
	Volume	Allows you to manage the lifecycle of vCenter volumes, such as creating, enabling, disabling, deleting, recovering, and expunging vCenter volumes.
		Allows you to attach/detach a volume to/from an instance.
		Allows you to change the owner of a vCenter volume.
	Image	Supported vCenter image types: system images in the VMDK format and volume images in VMDK format.
		Allows you to select the image platform. Supported platforms : Linux, Windows, and Other.
		Allows you to upload a vCenter image by using a URL.
		Allows you to manage the lifecycle of vCenter images, such as adding, enabling, disabling, deleting, recovering, and expunging vCenter images.
		Allows you to set the sharing mode for a vCenter image. Valid values: share globally, share to specified projects or accounts, and not share.
		Allows you to change the owner of a vCenter image.
	Event Message	Provides a list to centrally display event alarm messages of the vCenter, helping you locate problems quickly.
		Allows you to view event messages in a specified time period .
	Multi-account Management	Allows a tenant/sub-account to manage the lifecycle of resources such as VM instances and volumes of a vCenter it belongs to.
		Allows a tenant/sub-account to use vCenter resources such as networks and images shared by the admin.
		Allows a tenant/sub-account to view the usage of KVM VM instances and vCenter VM instances on the dashboard.
		Allows a tenant/sub-account to view the billing information of KVM and vCenter resources.
		Allows a tenant to apply for vCenter VM instances by submitting tickets.

Type	Features	Description
	Audit	Audits all of the vCenter actions, which effectively ensures the security of the cloud environment.

Features in Tenant Management

Type	Feature	Description
Personnel and Permissions	Organization	The basic element constructing organization structures. An organization structure consists of organizations of various levels.
		Provides a tree diagram to show the organizations in an organization structure. The admin or platform managers see all structure trees on the Cloud, while a normal platform or project member see only the tree its organization belongs to.
		Divides organizations into the default department and custom departments according to the users they organize. A custom department is used to organize personnel assigned to this department, and the default department is used to organize personnel has not been assigned to any custom department. Once a personnel is assigned to a custom department, it is removed from the default department.
		The default department is generated automatically by the system. You cannot delete the default department or add a sub-department to it.
		Allows you to centrally manage immediate members of the default department.
		Divides custom departments into two types according their addition methods: creating on local and synchronizing from a 3rd-party platform. The first method creates a custom department to organize local users and the second method provides a custom department to organize 3rd-party users.
		Divides custom departments into new teams and sub-departments according to their structural levels. A new team is a top-level department that allows you to add sub-departments of various level to it.
		Allows you to manage the lifecycle of a custom department, such as creating and deleting a custom department.
		Allows you to add sub-departments to a custom department or change the superior department for a sub-department.

Type	Feature	Description
		Allows you to set a department manager for a top-level department and department admins for custom departments.
		Allows you to centrally manage the immediate members and associated project resources of a custom department.
		Allows you to set quotas on custom department resources, such as the compute resource quota, storage resource quota , network resource quota, and other resource quota.
	User	Natural persons performing as the most basic units in Tenant Management.
		Divides users into local users and 3rd-party users according to their origins. Local users are created on the Cloud while 3rd-party users are synchronized from 3-party platforms.
		Allows you to manage the lifecycle of a local user, such as creating and deleting a local user.
		Supports two methods to create local users: custom creation and template import.
		Allows you to change the login password for a local user.
		Allows you to enable the certificate login feature for a local user to authenticate its identity when it logs in to the Cloud.
		Allows you to change a deleted AD/LDAP user from a 3rd-party user to a local user.
		Allows you to delete a 3rd-party user.
		Allows you to add/remove a user to/from a department, user group, or project.
		Allows you to set a platform or project role for a user.
		Allows you to specify a zone for a user to manage.
		Allows you to export the user information as a CSV table, which helps in statistical analysis and problem diagnosis.
	User Group	A collection of natural persons as well as a collection of project members.
		Allows you to manage the lifecycle of a user group, such as creating and deleting a user group.
		Allows you to add users to a user group and centrally manage the users in the user group.

Type	Feature	Description
		Allows you to add a user group to a project and assign unified project roles to the users in the user group.
	Role	A collection of permissions, granting users and user groups with permissions to perform actions on resources with APIs.
		Divides roles into platform roles and project roles according to the scenarios in which their permissions take effect. A platform role has permissions to manage the zone assigned to it while a project member has permissions to manage the project it belongs to.
		Divides roles into system roles and custom roles according to their generation mechanisms.
		System roles including admin, platform manager, department manager, monitor role, project admin, and project manager. Roles other than these are all custom roles.
		System roles are generated by the system automatically. You can view the UI permissions and API permissions of a system role.
		Allows you to manage the lifecycle of a custom role, such as creating and deleting a custom role.
		Allows you to modify the UI permissions and UI permissions of a custom role.
		Allows you to view the users and user groups bond with a role.
	3rd-party Authentication	Supports seamlessly access to 3rd-party authentication systems to enable 3rd-party users to log in to the Cloud.
		Supports 3rd-party authentication server types: AD, LDAP, OIDC, OAuth2, and CAS.
		Allows you to enable the SSL/TSL encryption for AD and LDAP servers.
		Allows you to enable SSL Certificate Check Skipping for LDAP servers configured with SSL certificates in Global Setting to skip all SSL certificate checks when the Cloud accesses these servers.
		Allows you to configure allowlist or blocklist filter mechanism and filter rules for an AD or LDAP server to filter the users

Type	Feature	Description
		that does not need or need to be synchronized from the base DN.
		Allows you to configure synchronize mapping rules for a 3rd-party authentication server.
		Allows you to manage the lifecycle of a 3rd-party authentication server, including adding and deleting a 3rd-party authentication server.
		Allows you to manually synchronize the latest user information from a AD or LDAP server.
		Allows you to manually test the connectivity of a AD or LDAP server.
	Audit	Audits all personnel and permissions actions, which effectively ensures the security of the cloud environment.
Project Management	Project	A project is a tenant. You can plan resources based on projects and create a separate resource pool for a project.
		Supports two project configuration methods: manual configuration and configuration with a project template.
		Supports project reclaim policies: unlimited, reclaim by specifying time, and reclaim by specifying cost.
		A project set as reclaimed by specifying time or reclaimed by specifying cost allows you to specify one of the following reclaim actions: disable project member login, disable project login and stop project resource, and delete project.
		Allows you to set an access control for a project as needed, allowing project members to log in to the Cloud during a specified time period, or prohibiting project member from logging in to the Cloud during a specified time period.
		Allows you to enable security group constraint for a project to associate a security group by force to each VM instance created by the project members.
		Allows you to manage the lifecycle of a project, such as creating, enabling, disabling, deleting, recovering, and expunging a project.
		Allows you to restore an expired project. The project member can log in to the project and the project resources work normally after the restoration.

Type	Feature	Description
		Allows you to generate a project template from an existing project for the fast creation of later projects.
		Allows you to set a project admin who can set project managers to help the project management.
		Allows you to set a department for a project. The project bill is merged into the department bill.
		Allows you to stop project resources, including VM instances and VPC vRouters. This action does not disable the project members from logging in to the project.
		Allows you to set quotas on project resources, such as compute resource quota, storage resource quota, network resource quota, and other resource quota.
		Allows you to centrally manage the members, user groups, associated resources, and shared resources of a project.
	Project Template	A template that identifies various resource quotas. You can use a project template to create a template quickly.
		Allows you to manage the lifecycle of a project template, such as creating and deleting a project template.
		Allows you to set quotas for a project template, such as compute resource quota, storage resource quota, network resource quota, and other resource quota.
	Audit	Audits all project management actions, which effectively ensures the security of the cloud environment.
Ticket Management	Process Management	Helps you provide basic resources to project more efficiently.
		Divides processes into the default process and custom processes according to their generation mechanisms.
		The default process is generated by the system and consists of two flows: the submitting ticket flow and the final approval and execution flow. The default process allows project admins, project managers, and normal project members to submit tickets, and the admin to approve and execute tickets.
		The default process applies to following tickets: tickets to modify project cycle, tickets to modify project quota, and tickets that are not specified with a custom process.
		Allows the admin, platform managers and normal platform members with corresponding permissions to create custom

Type	Feature	Description
		processes. A custom process consists of following flows: the submitting ticket flow, intermediate approval flows, and the final approval and execution flow. A custom process allows project admins, project managers, and normal project members to submit tickets; project admins, project mangers , normal project members, and department managers joining the projects to be responsible for intermediate approval flows , and the admin and project admins to be responsible for the final approval and execution flow.
		A custom process applies to following tickets: tickets to apply for VM instance, tickets to delete VM instance, and tickets to modify VM configuration.
		Supports the process lifecycle management, such as creating, enabling, disabling, and deleting a custom process.
		Supports ticket flow modifications.
	Ticket Application	Allows project admins, project managers, and normal project members to submit tickets.
		Supports following ticket types: apply for VM instance (KVM/ ESX), delete VM instance, modify VM configuration, modify project quota, and modify project cycle.
		Supports ticket lifecycle management, such as creating and deleting a ticket.
		Allows project members to recall a pending ticket, or resubmit a recalled or rejected ticket.
		Provides intuitive ticket processing records.
	Ticket Approval	Allows project admins, project managers, normal project members, and department managers joining the projects to be responsible for intermediate approval flows. Allows the admin and project admins to be responsible for the final approval and execution flow.
		Allows you to view pending and resolved tickets.
		Allows you to approve or reject a pending ticket.
		Provides intuitive ticket processing records.
		Allows the admin to view archived tickets, including resolved tickets that are deleted.

Type	Feature	Description
	Audit	Audits all ticket management actions, which effectively ensures the security of the cloud environment.

Features in Backup Service

Type	Features	Description
Backup Service	Backup Job Dashboard	Supports intuitive viewing and unified management of backup jobs on the Cloud to improve O&M efficiency.
		Displays backup job overview on different cards, including the number, state, and status of backup jobs.
		Displays backup job statistics in line chart and list format.
		Allows you to set a time filter to view the execution of backup jobs within the selected time period. The time filter applies to both the line chart and list.
		Allows you to view backup job execution details.
	Backup Job	Allows you to create a backup job to back up local VM instances, volumes, or databases to a specified storage server. Local backup, remote backup, and Public Cloud backup are currently supported.
		Allows you to manage the lifecycle of backup jobs, such as creating, enabling, disabling, and deleting backup jobs.
		Allows you to specify a local backup server for a backup job. If two local backup servers are specified, the failover mechanism is supported.
		Allows you to specify a remote backup server for a backup job. Only one remote backup server is supported. Supported types: Remote Backup and Alibaba Cloud Backup.
		Allows you to set a network QoS and disk QoS for VM/ volume backup jobs.
		Allows you to back up a VM instance with its attached volumes.
		Allows you to set a backup mode for a VM instance/volume backup job (incremental backup + default full backup, incremental backup + custom full backup).
		Allows you to specify a backup mode for a backup job of management node database (full backup mode).

Type	Features	Description
		Supports backup immediately after the job creation (VM instances/volumes backup jobs only).
		Allows you to manually perform a backup job, providing convenience for backing up important operations at any time.
		Allows you to set a data retention policy for a backup job, including local retention policy (by count/by time) and remote retention policy (permanently/by count/by time).
		Allows you to manage the backup resource of a backup job, including associating, disassociating, and viewing monitoring data in real time.
		Allows you to set a time filter to view backup job records within the selected time period.
		Significantly improves large file backup performance by optimizing the large file backup mechanism, supporting both physical and virtual tape libraries (requires tape library to provide file system mounting software, such as LTFS).
	Local Backup Data	Allows you to view the local backup data of VM instances, volumes, and databases in a list format.
		Allows you to view the backup data usage statistics of VM instances and volumes, including dependent incremental, incremental, and full.
		Allows you to recover the local backup data of VM instances/ volumes to local. Supported recovery policy: New Resource and Overwrite Original Resource.
		Allows you to recover a VM instance with its attached volumes. (The local backup data of the VM instance needs to contain volume backup data.)
		Allows you to recover local backup data from management node database to local.
		Allows you to change the owner of the local backup data of a VM instance.
		Allows you to scan a local backup server, and displays local backup data of the management node database on the cloud platform.

Type	Features	Description
		Allows you to export the local backup data of the management node database to the specified path of the local backup server, which is available for download.
		Allows you to delete the local backup data.
		Allows you to view the details of the VM/volume local backup data.
	Local Backup Server	Supports two types of addition: Existing Backup Storage (ImageStore only) and Add Server.
		Allows you to specify the backup network. In local backup scenarios, both data backup and recovery are implemented by using the backup network.
		Allows you to manage the lifecycle of local backup servers , such as creating, enabling, disabling, reconnecting, and deleting local backup servers.
		Allows you to scan a local backup server and display the local backup data record on the cloud platform.
		Allows you to clean up the invalid backup data and expired temporary data that have been completely deleted from the local backup server to free up the storage space.
		Allows you to update the password of the local backup server.
		Allows you to manage the backup data on local backup server, including VM instances, volumes, and the local backup data on database.
		Displays local backup server resource in a real time by using monitors, including capacity percent used, NIC, CPU, memory, and disk.
	Remote Backup Server	Allows you to add only one remote backup server. Supported types: Remote Backup and Alibaba Cloud Backup.
		Allows you to manage the lifecycle of remote backup servers , such as adding, enabling, disabling, reconnecting, and deleting remote backup servers.
		Allows you to update the password of a remote backup server.

Type	Features	Description
		Allows you to clean up the invalid backup data and expired temporary data that have been completely deleted from a remote backup server to free up the storage space.
		Allows you to manage the resources on a remote backup server, including backup data (VM instances, volumes, and the remote backup data on database) and zone.
	Remote Backup Data	Allows you to view the remote backup data of VM instances, volumes and management node database in a list format.
		Allows you to synchronize the remote backup data of VM instances and volumes to a local backup server.
		Allows you to recover the remote backup data of VM instances/volumes to local. Note that the remote backup data needs to synchronize to local backup server first before recovering to local.
		Allows you to recover the remote backup data of management node database to local.
		Allows you to scan the remote backup server, and display remote backup data of the management node database on the cloud platform.
		Allows you to export the remote backup data of the management node database to the specified path of the remote backup server, which is available for download.
		Allows you to delete the remote backup data.
	Audit	Audits all of the backup service actions, which effectively ensures the security of the cloud environment.

Features in Continuous Data Protection (CDP) Service

Type	Features	Description
Continuous Data Protection (CDP)	CDP Dashboard	Displays the critical CDP information on different cards, including the number and status of CDP tasks and recovery tasks, the CPU and memory utilization of backup servers, top 5 backup server usage, the total disk I/O of backup servers, and unread alarm statistics in recent 7 days.
	CDP Task	Allows you to create CDP tasks to continuously back up your VM data to a specified backup server to achieve continuous data protection.

Type	Features	Description
		Allows you to create CDP tasks in bulk for multiple VM instances. One CDP task corresponds to one VM instance.
		Allows you to perform a full backup for VM instances without installing any third-party agent.
		Performs a full backup for VM instances immediately after you create CDP tasks.
		Supports second/minute-level RPO settings
		Recommends the desired capacity required by a CDP task based on an algorithm when you create a CDP task for the first time, helping you to plan the backup space reasonably.
		Supports multiple primary storage: The CDP service applies to VM instances in different primary storage scenarios , including LocalStorage, NFS, SharedBlock, and Ceph primary storage.
		Allows you to manage the lifecycle of CDP tasks, such as creating, enabling, disabling, and deleting CDP tasks.
		Allows you to manually cancel the process of enabling a CDP task.
		Allows you to modify the protection policy of a disabled CDP task, including the recovery point interval, backup aggregation frequency, recovery point retention policy, and the backup rate.
		Allows you to modify the task running policy to adjust the desired size and RPO policy for a CDP task.
		Allows you to view the creation progress of a CDP task.
		Provides CDP task resource alarms and event alarms and allows you to create these alarms.
	CDP Data	Allows you to back up CDP data on a local backup server.
		Displays the CDP running status in charts and tables and allows you to view the details by specifying a time span.
		Displays hourly data changes so that you plan the backup capacity more reasonably.
		Provides a recovery point calendar, which identifies the dates with recovery points with colors and helps you to locate recovery points quickly.

Type	Features	Description
		Allows you to lock recovery points. After a recovery point is locked, data of the recovery point will not be automatically cleared or deleted.
		Provides recovery point list and locked recovery point list and allows you to view the details by specifying a time span.
		Supports fast recovery based on selected recovery points (including locked recovery points).
		Supports instant recovery with a minimum RTO in seconds.
		Supports entire restoration and file-level restoration.
		Entire restoration allows you to restore data to the original VM instance or to a newly-created VM instance.
		Restore data to a newly-created VM instance: Allows you to create a VM instance from the selected recovery point without affecting the original VM instance.
		The newly created VM instance will quickly start up for business recovery.
		Restore data to the original VM instance: Allows you to create new volumes or overwrite current volumes.
		Create new volumes: Allows you to retain and attach volumes before the recovery to the original VM instance to ensure data security.
		Overwrite current volumes: Overwrites the original data in the VM instance and retain the snapshots in the current volumes.
		After the data restoration, the original VM instance will quickly start up for business recovery.
		File-level restoration allows you to retrieve files without restoring the system. Supported file format include picture, text, and PDF.
		Allows you to clear CDP data, which will delete all the CDP data of the VM instance, including the locked recovery points . The Cloud performs full backup for the VM instance the next time the CDP task is enabled.
	Recovery Task	Provides a list of recovery tasks, allowing you to view the recovery records and progress for later audits and traceback.

Type	Features	Description
		Allows you to restore data through a wizard-style process.
		Supports multiple primary storage: The CDP service applies to VM instances in different primary storage scenarios, including LocalStorage, NFS, SharedBlock, and Ceph primary storage.
		Supports instant recovery with a minimum RTO in seconds.
		Allows you to restore data to the original VM instance or to a newly-created VM instance.
		Restore data to a newly-created VM instance: Allows you to create a VM instance from the selected recovery point without affecting the original VM instance.
		The newly created VM instance will quickly start up for business recovery.
		Restore to the original VM instance: Allows you to create new volumes or overwrite current volumes.
		Create new volumes: Allows you to retain and attach volumes before the recovery to the original VM instance to ensure data security
		Overwrite current volumes: Overwrites the original data in the VM instance and retain the snapshots in the current volumes.
		After the data restoration, the original VM instance will quickly start up for business recovery.
		Allows you to manage the lifecycle of recovery tasks, such as creating, enabling, disabling, and deleting recovery tasks.
		Allows you to redo a failed or canceled recovery task.
		Allows you to cancel a recovery task during the recovery progress. After a recovery task is canceled, intermediate data generated during the recovery process will not be retained.
	Local Backup Server	Allows you to back up CDP data on a local backup server.
		Allows you to use the ImageStore deployed in your local data center as the local backup server, or deploy a new local backup server.
		Allows you to add multiple local backup servers.

Type	Features	Description
		Allows you to view the CDP data saved to a local backup server on a local backup server details page.
	Audit	Audits all of the CDP actions, which effectively ensures the security of the cloud environment.

Features in Migration Service

Type	Features	Description
Migration Service	V2V Migration (VMware → the Cloud)	Allows you to migrate VM instances from a taken-over vCenter to the Cloud.
		Supported vCenter versions: 5.5, 6.0, 6.5, 6.7, and 7.0.
		Supported vCenter VM operating systems: RHEL 4.x/5.x/6.x/7.x, CentOS 4.x/5.x/6.x/7.x, SLES 11/12/15, Ubuntu 12/14/16/18, and Windows 7/Server 2003 R2/Server 2008 R2/Server 2012 R2/Server 2016/Server 2019.
		Supported source primary storage: Unlimited.
		Supported destination primary storage: LocalStorage, NFS, Ceph, and Shared Block.
		Allows you to manage the lifecycle of V2V jobs, including creating, rebooting, and deleting V2V jobs.
		Allows you to create V2V jobs for VM instances in bulk. The Cloud supports one V2V job per source VM instance.
		Allows you to enable the compression mode as needed, which effectively compresses the migration data cache and improves the cache space utilization of the V2V conversion host.
		Allows you to customize the configurations of destination VM instances.
		Allows you to view progress bars of V2V jobs.
		Automatically installs Windows VirtIO drivers for Windows VM instances during the migration process, which improves the NIC and disk operating efficiency.
	V2V Migration (KVM → the Cloud)	Allows you to migrate VM instances from a KVM platform to the Cloud.
		Allows you to migrate running or paused VM instances.

Type	Features	Description
		Supported source primary storage: Unlimited.
		Supported destination primary storage: LocalStorage, NFS, Ceph, and Shared Block.
		If the source primary storage or the destination primary storage is a Ceph storage, make sure that the libvirt is of 1.2 .16 or above version, and QEMU version is of 1.1 or above version before you perform the V2V migration. If neither the source primary storage nor the destination primary storage is a Ceph storage, make sure that the libvirt is of 1.2.9 or above version, and QEMU is of 1.1 or above version before you perform the V2V migration.
		Allows you to manage the lifecycle of V2V jobs, including creating, rebooting, and deleting V2V jobs.
		Allows you to create V2V jobs for VM instances in bulk. The Cloud supports one V2V job per source VM instance.
		Allows you to enable the compression mode as needed, which effectively compresses the migration data cache and improves the cache space utilization of the V2V conversion host.
		Allows you to customize the configurations of destination VM instances.
		Allows you to view progress bars of V2V jobs.
	V2V Conversion Host	Allows you to specify a host in the destination cluster as a V2V conversion host. The migration data is firstly cached in the V2V conversion host and then migrated to the destination primary storage.
		Allows you to attach data volumes to a V2V conversion host, so that you can cache data to your local disk or data volume as needed.
		Allows you to manage the lifecycle of V2V conversion hosts , such as adding, enabling, disabling, and deleting V2V conversion hosts.
		Make sure that the type of the V2V conversion host is consistent with that of the source platform.
		The state of a V2V conversion host is decoupled from that of the host added as the V2V conversion host. When the V2V conversion host is enabled but the host is disabled, the V2V

Type	Features	Description
		conversion host is used exclusively for V2V migrations, and other VM instances will not be dispatched to this host. This improves the migration efficiency.
		Allows you to set an independent migration network and network QoS to control transmission bottleneck and improve the migration efficiency.
		Monitors and displays the capacity usage of V2V conversion hosts.
	Audit	Audits all of the V2V actions, which effectively ensures the security of the cloud environment.

Features in Baremetal Management

Type	Features	Description
Baremetal Management	Baremetal Cluster	Provides independent cluster management for baremetal chassis.
		Allows you to manage the lifecycle of baremetal clusters, such as creating, enabling, disabling, and deleting baremetal clusters.
		Allows you to attach/detach a deployment server to/from a baremetal cluster.
		Allows you to attach/detach L2 networks to/from a baremetal cluster.
		Allows you to centrally manage the resources associated with a baremetal cluster, such as the deployment server, baremetal chassis, and L2 networks.
	Deployment Server	Allows you to specify an independent server as the deployment server to provide PXE services and console proxies for baremetal chassis.
		Allows you to manage the lifecycle of deployment servers , such as creating, enabling, disabling, reconnecting, and deleting deployment servers.
		Allows you to attach/detach baremetal clusters to/from a deployment server.
	Baremetal Chassis	Allows you to create baremetal instances based on baremetal chassis, which can be uniquely identified by their BMC interfaces and IPMI configurations.

Type	Features	Description
		Supports two types of addition: manual addition and template import. You can add up to 500 baremetal chassis at a time.
		Allows you to manage the lifecycle of baremetal chassis, such as adding, enabling, disabling, powering on, powering off, rebooting, and deleting baremetal chassis.
		Allows you to automatically or manually obtain the hardware information of a baremetal chassis.
		Allows you to launch the console of a baremetal chassis and jump to its IPMI management page.
		Allows you to view the hardware configuration of a baremetal chassis in a list format.
	Preconfigured Template	Quickly generates preconfigured files to achieve unattended bulk installation of baremetal instance operating systems.
		Divides preconfigured templates into system templates and custom templates based on how the preconfigured template is created.
		System templates are provided by the Cloud, which include the basic system variables and can be applied to simple unattended deployment scenarios.
		Custom templates are generated from the uploaded custom template files (in the UTF8 format), which include custom variables in addition to the basic system variables, and can be applied to complex unattended deployment scenarios.
		Supports the following operating systems: the custom OSs of the Cloud, mainstream Linux OSs (RHEL/CentOS series , Debian/Ubuntu series, and SUSE/openSUSE series), and other OSs.
		Supports the following types of template: kickstart (applies to the custom OSs of the Cloud, and RHEL/CentOS OSs), preseed (applies to Debian/Ubuntu OSs), and autoyast (applies to SUSE/openSUSE OSs).
		Allows you to manage the lifecycle of custom templates, such as adding, enabling, disabling, and deleting custom templates.
		Allows you to download a preconfigured template.
		Allows you to view the details of a preconfigured template.

Type	Features	Description
	Baremetal Instance	Created based on baremetal chassis as virtual instances of the baremetal chassis. You can add up to 50 baremetal instances at a time.
		Allows you to select images (in ISO format and are not live CDs) to deploy operating systems for baremetal instances.
		Allows you to achieve unattended bulk installation of baremetal instance operating systems with preconfigured files generated from the preconfigured templates.
		Allows you to configure business networks for a baremetal instance. Supports the following networks: flat network and public network. Supports the following network devices: NICs and NIC bonds.
		Allows you to manage the lifecycle of baremetal instances , such as creating, starting, stopping, rebooting, deleting, recovering, and expunging baremetal instances.
		Allows you to launch the console of a baremetal instance.
		Allows you to customize tags for baremetal instances so that you can locate them quickly.
		Supports internal monitoring: displays the baremetal instance data such as CPU, memory, disk I/O, disk size, and NIC I/O . An agent is required for internal monitoring.
		Allows you to centrally view the resources associated with a baremetal instance, such as NICs and disks.
	Audit	Audits all of the baremetal management actions, which effectively ensure the security of the cloud environment.

Features in Elastic Baremetal Management

Type	Resource	Description
Elastic Baremetal Management	Quick Start Wizard	Visualizes and displays the logical architecture of elastic baremetal management feature, guiding you to quickly use the elastic baremetal management.
		Provides five quick start steps, including Preparation → Provision Network → Elastic Baremetal Cluster → Gateway Node → Baremetal Node. After finishing the quick start wizard, you can create elastic baremetal instances. For NexaVM Ceph Enterprise, you need

Type	Resource	Description
		to make sure that the configuration is correct before creating elastic baremetal instances.
	Provision Network	Specifies a dedicated network for PXE processes and image downloading when elastic baremetal instances are created.
		Supported network type: IPv4.
		Allows you to manage the lifecycle of provision networks, such as creating and deleting provision networks.
		Allows you to view the associated elastic baremetal clusters.
	Elastic Baremetal Cluster	Provides independent cluster managements for baremetal nodes.
		Allows you to set the CPU architecture of an elastic baremetal cluster, including x86_64 and aarch64.
		Allows you to manage the lifecycle of elastic baremetal clusters, such as creating, enabling, disabling, and deleting elastic baremetal clusters.
		Allows you to attach/detach an L2 network of the NoVLAN/VLAN type to/from an elastic baremetal cluster.
		Allows you to change provision network for an elastic baremetal cluster.
		Allows you to attach/detach primary storage of the Ceph/ Shared Block type to/from an elastic baremetal cluster.
		Allows you to centrally manage resources associated with an elastic baremetal cluster, including gateway node, baremetal node, primary storage, iSCSI storage, and L2 network.
	Gateway Node	Forwards traffics of the Cloud and elastic baremetal instances.
		Allows you to manage the lifecycle of gateway nodes, such as adding, enabling, reconnecting, and deleting gateway nodes.
		Allows you to change the password of a gateway node.
		Allows you to change elastic baremetal cluster of a gateway node.

Type	Resource	Description
		Monitors and displays gateway node metrics such as NIC , CPU, and memory.
		Allows you to centrally manage elastic baremetal instances associated with a gateway node.
	Baremetal Node	A baremetal node is used to create elastic baremetal instances and is universally identified by its BMC interface and IPMI configurations.
		Supports two types of addition: custom and template import. You can add up to 500 baremetal nodes at a time . (You can modify the maximum number of bulk addition in global setting.)
		Allows you to set the CPU architecture of a baremetal node, including x86_64 and aarch64.
		Allows you to set the start method of a baremetal node, including volume and local disk (non take-over/take-over).
		Allows you to manage the lifecycle of baremetal nodes , such as adding, enabling, disabling, powering on, powering off, rebooting, and deleting baremetal nodes.
		Allows you to automatically or manually obtain the hardware information of a baremetal node.
		Allows you to modify the IPMI info when the power supply of the baremetal node is in Unknown state.
		Allows you to launch the console of a baremetal node and jump to its IPMI management page.
		Allows you to view the hardware information of baremetal nodes in a list format.
	Elastic Baremetal Offering	An elastic baremetal offering defines the number of vCPU cores, memory size, CPU architecture, CPU model , and other configuration settings of elastic baremetal instances. You can use an elastic baremetal offering to create an elastic baremetal instance.
		Allows you to obtain an elastic baremetal offering by obtaining the hardware information of baremetal nodes.

Type	Resource	Description
		Allows you manage the lifecycle of elastic baremetal offerings, such as enabling and disabling elastic baremetal offerings.
		Allows you to set the sharing mode of an elastic baremetal offering, including share globally, share to specified projects or accounts, and not share.
		Allows you to centrally manage the baremetal nodes associated with an elastic baremetal offering.
	Elastic Baremetal Instance	Comparable to instances virtualized through physical servers in performance, leverages resource scalability in the Cloud to achieve flexible applications and on-demand usages.
		Supports two types of creation: add by baremetal node and add by baremetal offering.
		Allows you to power off to release baremetal node. When elastic baremetal instances are stopped, baremetal nodes will be automatically released to avoid idle resources (only elastic baremetal instances added by elastic baremetal offerings and baremetal nodes that start on volume).
		Allows you to specify the storage allocation policy of an elastic baremetal instance, including system allocation and custom (only elastic baremetal instances added by elastic baremetal offerings and baremetal nodes that start on volume).
		Allows you to select an image to install the operating system for an elastic baremetal instance. Supported operating systems: x86 Windows (2012/2016/2019/10), x86 Linux (CentOS 7/8, Ubuntu 18 LTS/20 LTS), and ARM Linux (CentOS 7/Kylin V10) (only elastic baremetal instances added by elastic baremetal offerings and baremetal nodes that start on volume/non take-over local disk).
		Allows you to specify the gateway node allocation policy, including LeastBmPreferredGatewayAllocationStrategy, Last Gateway Node, and Random. You can select a gateway node as the first assigned gateway node for an elastic baremetal instance.

Type	Resource	Description
		Allows you to manage the lifecycle of elastic baremetal instances, such as creating, starting, stopping, rebooting , powering off, deleting, recovering, and expunging elastic baremetal instances.
		Allows you to automatically or manually obtain the status of an elastic baremetal instance.
		Allows you to launch the console of a running elastic baremetal instance (agent required).
		Allows you to customize tags for elastic baremetal instances so that you can locate them quickly.
		Allows you to attach/detach a volume to/from an elastic baremetal instance (agent required).
		Allows you to attach/detach a block storage volume to/ from an elastic baremetal instance (agent required).
		Allows you to change system of an elastic baremetal instance.
		Allows you to change the password of an elastic baremetal instance (agent required).
		Allows you to create an image for an elastic baremetal instance (only elastic baremetal instances that start on volume).
		Allows you to create a single snapshot for an elastic baremetal instance (only elastic baremetal instances that start on volume).
		Monitors and displays elastic baremetal instance metrics such as CPU, memory, disk, disk capacity, and NIC (agent required).
		Allows you to configure business networks for elastic baremetal instances. Supported business network: flat network, public network, and VPC network. Supported network device: NIC and NIC Bond.
		Allows you to centrally manage resources associated with elastic baremetal instance, including volume, NIC (provision NIC and business NIC), and local disk.

Type	Resource	Description
	Audit	Audits all of the elastic baremetal management actions , which effectively ensure the security of the cloud environment.

Features in Hybrid Cloud Management

Type	Features	Description
Hybrid Cloud Management	Sync Data	Allows you to synchronize Alibaba Cloud resources from added regions and zones to local, such as ECS instances , disks, VPCs, vSwitches, security groups, images, EIPs, VPNs, virtual border routers, and router interfaces.
		Supports automatic and manual data synchronizations. Automatic synchronizations occur when regions or zones are newly added to local.
	Quick Start Wizard	Visualizes the logical architecture of Hybrid Cloud Management, guiding you to use Hybrid Cloud Management quickly.
		Provides three quick start steps: Create ECS Instance, Establish VPN Connection, and Create Alibaba Cloud Express Connect.
	ECS Instance	ECS is Elastic Compute Service provided by Alibaba Cloud.
		Allows you to manage the lifcycle of an ECS instance, such as creating, starting, stopping, rebooting, and deleting an ECS instance.
		Allows you to launch the console of an ECS instance and modify the console password as needed.
		Allows you to modify the system user password of an ECS instance. The new password takes effect after you reboot the ECS instance.
		Allows you to centrally manage disks attached to an ECS instance.
	Disk	Alibaba Cloud disks that provide extended storage spaces for ECS instances.
		Supports two types of disks: ultra cloud disks and SSD disks.
		Allows you to manage the lifecycle of a disk, such as creating and deleting a disk.

Type	Features	Description
		Allows you to attach/detach disks to/from ECS instances.
		Allows you to set whether to delete a disk simultaneously when you delete the ECS instance it attached to.
	Image	Alibaba Cloud images that provide template files to create ECS instances.
		Divides images into two types according to their origins: Alibaba Cloud images and custom images. Alibaba Cloud images are synchronized from Alibaba Cloud to local. Custom images are created locally and uploaded to Alibaba Cloud through buckets in corresponding regions.
		Allows you to choose the format of uploaded local images in Hybrid Cloud Settings. Valid values: .qcow2 and .raw.
		Displays the upload progress of local images.
		Allows you to delete images.
	VPC	Provides 3 CIDRs for you to create VPCs (Alibaba virtual private clouds) dedicated for ECS instances: 192.168.0.0/16 , 172.160.0.0/12, and 10.0.0.0/8
		Allows you to manage the lifecycle of a VPC, such as creating and deleting a VPC.
		Allows you to create VPN connections and express connects based on VPCs.
		Allows you to centrally manage associated resources of a VPC, such as vSwitches, vRouters, security groups, and VPN gateways.
		Allows you to manage the lifecycle of a vSwitch, such as creating and deleting a vSwitch.
		Allows you centrally manage the ECS instances associated with a vSwitch.
		Allows you to add/delete route entries to/from vRouters.
		Provides three next hop options for route entries: hop to route interface, hop to ECS instance, and hop to VPN gateway.
	VPN	An IPsec tunnel created between a VPN gateway and a VPN customer gateway that enables communications between local private networks and VPC networks on Alibaba Cloud.

Type	Features	Description
		A VPN gateway is a network connection service provided by Alibaba Cloud. You need to purchase it on Alibaba Cloud Console before you can use it.
		Allows you to delete a VPN gateway from local without influencing the corresponding actual resource on Alibaba Cloud.
		Allows you to centrally manage the VPN connections based on a VPN gateway.
		A VPN customer gateway provides services for the local data center.
		Allows you to manage the lifecycle of a VPN customer gateway, such as creating and deleting a VPN customer gateway.
		Allows you to centrally manage the VPN connections based on a VPN customer gateway.
		Allows you to establish a VPN connection between a VPN gateway and a VPN customer gateway to enable encrypted communications between the local data center and Alibaba Cloud.
		Provides three entries for you to create VPN connections: from Quick Start Wizard, from a VPC action list, and on the **VPN Connection** page.
		Allows you attach multiple local VPC networks to a VPN connection.
		Supports NAT Traversal that ensures normal data transmissions even though a NAT device exists between the local data center and Alibaba Cloud.
	Express Connect	A physical circuit to connect the local data center and the access point of Alibaba Cloud that ensures fast, stable and secure communications between local private networks and Alibaba Cloud VPCs.
		Provides 2 creation entries for express connect: from Quick Start Wizard and on a VPC action list.
		Allows you to centrally manage resources used for express connects, such as router interfaces and virtual border routers.

Type	Features	Description
		Allows you to add router interfaces to virtual border routers and VPC vRouters for message forwards.
		Allows you to specify regions to synchronize virtual border routers to local.
		Allows you to add/delete route entries to/from a virtual border router.
		Provides four next hop options for route entries: hop to ECS instance, hop to router interface, hop to VPN gateway, and hop to physical connection interface.
		Allows you to modify the interconnection address of a virtual border router.
	Security Group	Alibaba Cloud security groups that provide security control services for ECS instances on the L3 network layer.
		Provides four initial rule options for security groups: Prohibit All (Default), Allow All, Disable Some Vulnerable Ports, and Allow Commonly Used Ports.
		Allows you to manage the lifecycle of a security group, such as creating and deleting a security group.
		Allows you to add/delete ingress or egress rules to/from a security group.
		Provides two authorization policy options for ingress/egress rules: Accept and Reject.
		Provides five protocol options for ingress/egress rules: ALL, TCP, UDP, ICMP, and GRE.
		Allows you to set priorities for ingress/egress rules. The rule with the highest priority takes effect when you set multiple rules on a same object.
	EIP	Elastic IP addresses (EIP) in Alibaba Cloud public networks that enable ECS instances to access public networks.
		Allows you manage the lifecycle of an EIP, such as creating and deleting an EIP.
		Allows you to attach/detach EIPs to/from ECS instances.
	Alibaba Cloud NAS	Integrates Alibaba Cloud NAS to provide file systems as backend storage systems for AliyunNAS primary storage.

Type	Features	Description
		Supports two methods to add NAS file systems: add an existing file system deployed on Alibaba Cloud, or create a new file system.
		NAS files systems supports two storage types: Performance and Capacity.
		NAS file systems supports two protocol types: NFS and SMB .
		Allows you to manage the lifecycle of an NAS file system, such as creating and deleting a file system.
		Allows you to create permission groups to limit accesses to a file systems.
		Permission groups support allowlist mechanisms, allowing you to add rules to allow specified IP addresses and CIDRs to access the file system.
		Supports two methods to create permission groups: add an existing permission group on Alibaba Cloud, or create a new permission group.
		Allows you to add/delete rules to/from a permission group.
		Allows you to set the permission range when you create a permission group rule, enabling an authentication objects to only read from the file system (RDONLY), or read from as well as write in the file system (RDWR).
		Allows you to set priorities for permission group rules. The rule with the highest priority takes effect when you set multiple rules on a same authentication object.
		Allows you to create an AliyunNAS primary storage based on a file system and permission groups.
		AliyunNAS primary storage supports backup storage: ImageStorage backup storage.
		Allows you to manage the lifecycle of an AliyunNAS primary storage, such as adding, enabling, disabling, reconnecting, deleting an AliyunNAS primary storage or making it enter the maintenance mode.
		Allows you to centrally manage the resources associated with an AliyunNAS primary storage, such as VM instance, volumes, and clusters.

Type	Features	Description
		Allows you to clean up garbage data of an AliyunNAS at a specified interval. You can modify the interval in Hybrid Cloud Settings.
		Monitors and displays the percentage of used capacity of an AliyunNAS primary storage.
	Alibaba Cloud EBS	Integrates Alibaba Cloud EBS to serve as a local primary storage type, AliyunEBS.
		AliyunEBS primary storage supports backup storage: AliyunEBS backup storage.
		Allows you to manage the lifecycle of an AliyunEBS primary storage, such as adding, enabling, disabling, reconnecting, and deleting an AliyunEBS primary storage or making it enter the maintenance mode.
		Allows you to centrally manage the resources associated with an AliyunEBS primary storage, such as VM instances, volumes, and clusters.
		Monitors and displays the percentage of used capacity of an AliyunEBS primary storage.
		Allows you to clean up garbage data of an AliyunEBS at a specified interval. You can modify the interval in Hybrid Cloud Settings.
		Integrates Alibaba Cloud Object Storage Service (OSS) to serve as a local backup storage type, AliyunEBS.
		AliyunEBS backup storage supports primary storage: AliyunEBS primary storage.
		Allows you to set a dedicated data network for an AliyunEBS backup storage to improve the data transmission efficiency between compute nodes and the backup storage.
		Allows you to manage the lifecycle of an AliyunEBS backup storage, such as adding, enabling, disabling, reconnecting, and deleting an AliyunEBS backup storage.
		Allows you to centrally manage the images in an AliyunEBS backup storage.
		Monitors and displays the percentage of used capacity of an AliyunEBS backup storage.

Type	Features	Description
	Region	Allows you to add Alibaba Cloud regions can be accessed by your AccessKey. The zones and resources in the regions can be synchronized to local.
		Supports two types of regions: Alibaba Cloud regions and Private Alibaba Cloud regions.
		Divides Private Alibaba Cloud regions into two types: AliyunNAS region and AliyunEBS region.
		Allows you to centrally manage the zones and buckets in a region.
		Allows you to use a bucket to transfer a local image to Alibaba Cloud.
		Supports two methods to add buckets: add an available bucket existing in the region, or create a new bucket.
		Allows you to manage the lifecycle of a bucket, such as adding and deleting a bucket.
		Allows you to set a bucket as the default bucket for the image upload. Allows you to cancel the default state of a bucket.
		Allows you to delete a region from local without influencing the corresponding actual resource on Alibaba Cloud.
	Zone	Allows you to synchronize zones in a region you added , or manually add zones that can be accessed by your AccessKey. Resources in an added zone can be synchroniz ed to local.
		Allows you to centrally manage the resources associated with a zone, such as vSwithes and ECS instances.
		Allows you to delete a zone from local without influencing the corresponding actual resource on Alibaba Cloud.
	AccessKey Management	An identity credential that has access to APIs of Alibaba Cloud or Private Alibaba Cloud, thus enabling you to use relevant Cloud services.
		Supports two types of AccessKeys: Alibaba Cloud AccessKeys and Private Alibaba Cloud AccessKeys.
		Divides Private Alibaba Cloud AccessKeys into two types: AliyunNAS AccessKey and AliyunEBS AccessKey.

Type	Features	Description
		Allows you to manage the lifecycle of an AccessKey, such as adding and deleting an AccessKey.
		Allows you to set an AccessKey as default to call APIs of Alibaba Cloud or Private Alibaba Cloud. Allows you to cancel the default state of an AccessKey.
		Displays the basic information of an AccessKey, which helps in the user management.
	Hybrid Cloud Settings	Allows you to configure settings that take effect on the whole platform.
		Supports quick search and directory navigation to help you quickly locate target items.
		Allows you to modify settings of a single item in Hybrid Cloud Settings.
	Audit	Audits all of the Hybrid Cloud Management actions, which effectively ensure the security of the cloud environment.

4 Product Highlights

NexaVM Cloud is the next-generation software defined featuring Simple, Strong, Scalable and Smart (4S).

1. **Simple**

   - Easy installation and deployment: Provides installation packages on our official website. You can install and deploy the Cloud from scratch within just 30 minutes.
   - Easy to set up: Supports bulk VM operations, such as creating or deleting VM instances in bulk.
   - Simple, practical operations: Provides a thorough User Guide with ample help information, productive community, and standard APIs.
   - Friendly UI: Provides a well-designed user interface with powerful features at your fingertip.

2. **Strong**

   - Stable, efficient system architecture design: Provides an asynchronous architecture, in-process microservices architecture, lock-free architecture, stateless service architecture, and consistent hashing ring to ensure the system efficiency and stability. A single management node can manage tens of thousands hosts, and hundreds of thousands of VM instances. A cluster that contains multiple management nodes can use a database and a set of message buses to manage hundreds of thousands of hosts and millions of VM instances, and handle tens of thousands of concurrent APIs.
   - High concurrent API requests: A single NexaVM Cloud management node can easily handle tens of thousands of concurrent API call requests per second.
   - Stringent HA requirements: When a network or management node is unavailable, VM instances can be automatically switched to another management node that is detected as healthy. The management node virtualization helps to achieve the high availability for a single management node. That is, standby management nodes will be dynamically applied within seconds if any management node is disconnected, thus ensuring your business continuity.

3. **Scalable**

   - Large scale: A single management node can manage one to tens of thousands of hosts and hundreds of thousands of VM instances.

- Comprehensive API: NexaVM Cloud provides a whole set of IaaS APIs. Hence, you can create brand-new, available zones across multiple geographical locations, modify network configurations, and upgrade physical servers.
- Resource allocation based on your needs: Resizes important resources such as VM instances and cloud storages according to your demands. NexaVM Cloud not only allows you to modify online the CPU, memory, and other resources for a VM instance, but also allows you to dynamically adjust its network bandwidth, disk bandwidth, and other resources for a VM instance.

4. Smart

- Automatic O&M: Everything in NexaVM Cloud is managed APIs. By using the Ansible inventory, NexaVM Cloud can realize full-automatic deployment and upgrade as well as automatic detection and reconnection. If network jitters happen or hosts restart, each management node can be automatically reconnected to the networks or the hosts. Note that a NexaVM Cloud scheduler allows you to start or stop VM instances on schedule, and allows you to take VM snapshots on schedule with the round-robin policy.
- Online seamless upgrade: Provides one-click seamless upgrade within 5 minutes. Hence, you only need to upgrade and manipulate management nodes. After the Cloud is upgraded successfully and started, the compute node, storage node, and network node will be automatically upgraded as well.
- Real-time global monitoring: Manages and controls the current resource consumption of the entire cloud. With the real-time monitoring, you can adjust your resources intelligently to save IT software and hardware resources.

Glossary

Instance

An instance is a virtual machine or server that runs the images of operating systems in Cloud, such as VM instance and elastic baremetal instance.

VM Instance

A VM instance is a virtual machine instance running on a host. A VM instance has its own IP address and can access public networks and run application services.

Volume

A volume provides storage space for a VM instance. Volumes are categorized into root volumes and data volumes.

Root Volume

A root volume provides support for the system operations of a VM instance.

Data Volume

A data volume provides extended storage space for a VM instance.

Image

An image is a template file used to create a VM instance or volume. Images are categorized into system images and volume images.

Instance Offering

An instance offering defines the number of vCPU cores, memory size, network bandwidth, and other configuration settings of VM instances.

Disk Offering

A disk offering defines the capacity and other configuration settings of volumes.

GPU Specification

A GPU specification defines the frame per second (FPS), video memory, resolution, and other configuration settings of a physical or virtual GPU. GPU specifications are categorized into physical GPU specifications and virtual GPU specifications.

vNUMA Configuration

vNUMA uses CPU pinning to passthrough the topology of associated host physical NUMA (pNUMA) nodes to a VM instance, generating a topology of virtual NUMA (vNUMA) nodes for the VM instance. 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 VM instance.

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 VM instance 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 VM instance to physical CPUs (pCPUs) of the host so that these threads run on assigned pCPUs.

Auto-Scaling Group

An auto-scaling group is a group of VM instances that are used for the same scenarios. An auto-scaling group can automatically scale out or in based on application workloads or health status of VM instances in the group.

Snapshot

A snapshot is a point-in-time capture of data status in a volume.

Affinity Group

A VM scheduling policy is a resource orchestration policy based on which VM instances are assigned hosts to achieve the high performance and high availability of businesses.

Zone

A zone is a logical group of resources such as clusters, L2 networks, and primary storage. Zone is the largest resource scope defined in the Cloud.

Cluster

A cluster is a logical group of hosts (compute nodes).

Host

A host provides compute, network, and storage resources for VM instances.

Primary Storage

A primary storage is one or more servers that store volume files of VM instances. These files include root volume snapshots, data volume snapshots, image caches, root volumes, and data volumes.

Backup Storage

A backup storage is a storage server that stores VM image templates, including ISO image files.

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 Shared Block primary storage or pass through the block to a VM instance.

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 Shared Block primary storage or pass through the block to a VM instance.

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 SharedBlock primary storage.

L2 Network

An L2 network is a layer 2 broadcast domain used for layer 2 isolation. Generally, L2 networks are identified by names of devices on the physical network.

VXLAN Pool

A VXLAN pool is a collection of VXLAN networks established based on VXLAN Tunnel Endpoints (VTEPs). The VNI of each VXLAN network in a VXLAN pool must be unique.

L3 Network

An L3 network includes IP ranges, gateway, DNS, and other network configurations that are used by VM instances.

Public Network

Generally, a public network is a logical network that is connected to the Internet. However, in an environment that has no access to the Internet, you can also create a public network.

Flat Network

A flat network is connected to the network where the host is located and has direct access to the Internet. VM instances in a flat network can access public networks by using elastic IP addresses.

VPC Network

A VPC network is a private network where VM instances can be created. A VM instance in a VPC network can access the Internet through a VPC vRouter.

Management Network

A management network is used to manage physical resources in the Cloud. For example, you can create a management network to manage access to hosts, primary storages, backup storages, and VPC vRouters.

Flow Network

A flow network is a dedicated network for port mirror transmission. You can use a flow network to transmit the mirrors of data packets of NIC ports to the target ports.

VPC vRouter

A VPC vRouter is a dedicated VM instance that provides multiple network services.

VPC vRouter HA Group

A VPC vRouter HA group consists of two VPC vRouters. Either VPC vRouter can be a primary or secondary VPC vRouter for the group. If the primary VPC vRouter does not work as expected, the VPC vRouter becomes the secondary VPC vRouter in the group to ensure high availability of business.

vRouter Image

A vRouter image encapsulates network services and can be used to create VPC vRouters and load balancers. vRouter images can be categorized into VPC vRouter images and load balancer (LB) images.

Dedicated-Performance LB Image

A dedicated-performance load balancer (LB) image encapsulates dedicated-performance load -balancing services and can be used to create load balancer instances. However, a dedicated-performance load balancer image cannot be used to create VM instances.

vRouter Offering

A vRouter offering defines the number of vCPU cores, memory size, image, management network , and public network configuration settings of VPC vRouters. You can use a vRouter offering to create VPC vRouters that can provide network services for public networks and VPC networks.

LB Instance Offering

A load balancer (LB) instance offering defines the CPU, memory, image, and management network configuration settings used to create LB instances. LB instances provide load balancing services for the public network, flat network, and VPC network.

SDN Controller

An SDN controller is used to control network devices such as switches. You can add an external SDN controller to the Cloud and use the controller to control external switches and other network devices.

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.

VIP

In bridged network environments, a virtual IP address (VIP) provides network services such as serving as an elastic IP address (EIP), port forwarding, load balancing, IPsec tunneling. When a VIP provides the preceding network services, packets are sent to the VIP and then routed to the destination network where VM instances are located.

EIP

An elastic IP address (EIP) functions based on the NAT technology. IP addresses in a private network are translated into an EIP that is in another network. This way, private networks can be accessed from other networks by using EIPs.

Port Forwarding

Port forwarding functions based on the layer-3 forwarding service of VPC vRouters. This service forwards traffic flows of the specified IP addresses and ports in a public network to specified ports of VM instances by using the specified protocol. If your public IP addresses are insufficient, you can configure port forwarding for multiple VM instances by using one public IP address and port.

Load Balancer

A load balancer distributes traffic flows of a virtual IP address to backend servers. It automatically inspects the availability of backend servers and isolates unavailable servers during traffic distributi on. This way, the load balancer improves the availability and service capability of your business.

Listener

A listener monitors the frontend requests of a load balancer and distributes the requests to a backend server based on the specified policy. In addition, the listener performs health checks on backend servers.

Forwarding Rule

A forwarding rule forwards the requests from different domain names or URLs to different backend server groups.

Backend Server Group

A backend server group is a group of backend servers that handles requests distributed by load balancers. It is the basic unit for traffic distribution by load balancer instances.

Backend Server

A backend server handles requests distributed by a load balancer. You can add a VM instance on the Cloud or a server on a third-party cloud as a backend server.

Frontend Network

A frontend network is a type of network that is associated with a load balancer. Requests from the network are distributed by the load balancer to backend servers based on a specified policy.

Backend Network

A backend network is a type of network that is associated with a load balancer. Requests from frontend networks are distributed by the load balancer to servers in the backend network.

Load Balancer Instance

A load balancer instance is a custom VM instance used to provide load balancing services.

Certificate

If you select HTTPS for a listener, associate it with a certificate to make the listener take effect. You can upload either a certificate or certificate chain.

Firewall

A firewall is an access control policy that monitors ingress and egress traffic of VPC vRouters and decides whether to allow or block specific traffic based on the associated rule sets and rules.

Firewall Rule Set

A firewall rule set is a set of rules that a firewall uses to defend against network attacks. You need to associate a rule set with the egress or ingress flow direction of VPC vRouter NICs to make the rule set take effect.

Firewall Rule

A firewall rule is an access control entry associated with the egress or ingress flow direction of VPC vRouter NICs to defend against network attacks. A firewall rule includes rule priority, match condition, and behavior.

Rule Template

A rule template is a template that you can select when you add rules to a rule set or a firewall.

IP/Port Set

An IP or port set is a set of IP addresses or ports that you can select when you add rules to a rule set or a firewall.

IPsec Tunnel

An IPSec tunnel encrypts and verifies IP packets that transmit over a virtual private network (VPN) from one site to another.

OSPF Area

An Open Shortest Path First (OSPF) area is divided from an autonomous system based on the OSPF protocol. This simplifies the hierarchical management of vRouters.

NetFlow

A NetFlow monitors the ingress and egress traffic of the NICs of VPC vRouters. The supported versions of data flows are V5 and V9.

Port Mirroring

Port mirroring mirrors the traffic data of VM NICs and sends the traffic data to the target ports. This allows for the analysis of data packets of ports and simplifies the monitoring and management of data traffic and makes it easier to locate network errors and exceptions.

Route Table

A route table contains information about various routes that you configure. Route entries in a route table must include the destination network, next hop, and route priority.

CloudFormation

CloudFormation is a service that simplifies the management of cloud resources and automates deployment and O&S. You can create a stack template to configure cloud resources and their dependencies. This way, resources can be automatically configured and deployed in batches. CloudFormation provides easy management of the lifecycle of cloud resources and integrates automatic O&S into API and SDK.

Resource Stack

A resource stack is a stack of resources that are configured by using a stack template. The resources in the stack have dependencies with each other. You can manage resources in the stack by managing the resource stack.

Stack Template

A stack template is a UTF8-encoded file based on which you can create resource stacks. The stack template defines the resources that you want, the dependencies between the resources , and the configuration settings of the resources. When you use a stack template to create a resource stack, CloudFormation parses the template and the resources are automatically created and configured.

Sample Template

A sample template is a commonly used resource stack. You can use a sample template provide by the Cloud to create resource stacks.

Designer

A designer is a CloudFormation tool that allows you to orchestrate cloud resources. You can drag and drop resources on a canvas and use lines to establish dependencies between the resources.

Baremetal Cluster

A baremetal cluster consists of baremetal chassis. You can manage baremetal chassis by managing a baremetal cluster where the chassis reside.

Deployment Server

A deployment server is a server that provides PXE service and console proxy service for baremetal chassis.

Baremetal Chassis

A baremetal chassis is used to create a baremetal instance and is identified based on the BMC interface and IPMI configuration setting.

Preconfigured Template

A preconfigured template is used to create a preconfigured file that allows for unattended batch installation of an operating system for baremetal instances.

Baremetal Instance

A baremetal instance is an instantiated baremetal chassis.

Elastic Baremetal Management

Elastic Baremetal Management provides dedicated physical servers for your applications to ensure high performance and stability. In addition, this feature allows elastic scaling. You can apply for and scale resources based on your needs.

Provision Network

A provision network is a dedicated network for PXE boot and image downloads while creating elastic baremetal instances.

Elastic Baremetal Cluster

An elastic baremetal cluster consists of elastic baremetal instances. You can manage elastic baremetal instances by managing an elastic baremetal cluster where the instances reside.

Gateway Node

A gateway node is a node where the ingress and egress traffic of the Cloud and elastic baremetal instances is forwarded.

Baremetal Node

A baremetal node is used to create a baremetal instance and is identified based on the BMC interface and IPMI configuration setting.

Elastic Baremetal Instance

An elastic baremetal instance has the same performance as physical servers and allows elastic scaling. You can apply for and scale resources based on your needs.

Elastic Baremetal Offering

An elastic baremetal offering defines the number of vCPU cores, memory size, CPU architecture, CPU model, and other configuration settings of elastic baremetal instances.

vCenter

The Cloud allows you to take over vCenter and manage resources on the vCenter.

VM Instance

A VM instance is an ESXi virtual machine instance running on a host. A VM instance has its own IP address to access public networks and can run application services.

Network

A vCenter network defines the network settings of VM instances on vCenter, such as IP range, gateway, DNS, and network services.

Volume

A volume provides storage space for a VM instance on vCenter. A volume attached to a VM instance can be used as a root volume or data volume. A root volume provides support for the system operations of a VM instance. A data volume provides extended storage space for a VM instance.

Image

An image is a template file used to create a VM instance or volume on vCenter. Images are categorized into system images and volume images.

Event Message

Event Message displays event alarm messages of vCenter that is took over by the Cloud. This feature allows you to locate errors and exceptions efficiently.

Network Topology

A network topology visualizes the network architecture of the Cloud. It allows for efficient planning , management, and improvement of network architecture. Network topologies can be categorized into global topologies and custom topologies.

Performance Analysis

Performance Analysis displays the performance metrics of key resources monitored externally or internally in the Cloud. You can view the performance analysis or export the analysis report as needed to improve the O&M efficiency.

Capacity Management

Capacity Management visualizes the capacities and usages of key resources in the Cloud. You can use this feature to improve O&S efficiency.

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, event alarm, and extended alarm.

One-Click Alarm

A one-click alarm integrates multiple metrics of a resource. You can create one-click alarms for multiple resources to monitor these resources.

Alarm Template

An alarm template is a template of alarm rules. If you associate an alarm template with a resource group, an alarm is created to monitor the resources in the group.

Resource Group

A resource group consists of resources grouped based on your business needs. If you associate an alarm template with a resource group, the alarm rules specified by the template take effect on all the resources in the group.

Message Template

A message template specifies the text template of a resource alarm message or event alarm message sent to an SNS system.

Message Source

A message source is used to take over extended alarm messages. If you configure alarms for message sources, extended alarm messages can be sent to various endpoints.

Endpoint

An endpoint is a method that users obtain subscribed messages. Endpoints are categorized into system endpoints, email, DingTalk, HTTP application, short message service, and Microsoft Teams.

Alarm Message

An alarm message is a message sent the time when an alarm is triggered.

Current Task

A current task is an ongoing operation performed in the Cloud. You can perform centralized management over ongoing operations.

Operation Log

An operation log is a chronological record of operations on the specified objects and their operation results.

Audit

Audit monitors and records all activities on the Cloud. You can use this feature to implement operation tracking, cybersecurity classified protection compliance, security analysis, troublesho oting, and automatic O&M.

Log Collection

Allows you to collect with one click the log data from the Cloud and various nodes on the Cloud generated in the specified time period and download the log data.

One-Click Inspection

Comprehensively inspects the health status of key resources and services of the Cloud and scores their healthiness based on the inspection results. In addition, the one-click inspection service provides O&M suggestions and inspection reports.

Backup Management

Backup management integrates multiple disaster recovery technologies such as incremental backup and full backup that are suitable for multiple business scenarios. You can implement local backup and remote backup based on your business needs.

Backup Job

You can create a backup job to back up local VM instances, volumes, or databases to a specified storage server on a regular basis.

Local Backup Data

Local backup data of VM instances, volumes, and databases is stored in the local backup storage.

Local Backup Server

A local backup server is located at the local data center and is used to store local backup data.

Remote Backup Server

A remote backup server is located at a remote data center or a public cloud and is used to store remote backup data.

Continuous Data Protection (CDP)

Continuous Data Protection (CDP) provides second-level and fine-grained continuous backups for important business systems in VM instances, allowing users to restore VM data to a specific time state, and retrieve files without restoring the system.

CDP Task

You can create a CDP task to continuously back up your VM data to a specified backup server to achieve continuous data protection and recovery.

CDP Data

The backup data generated from continuous data protection on VM instances is stored in local backup servers.

Recovery Point

A recovery point is a data point generated during continuous data protection. A recovery point corresponds to a data record within the recovery point interval specified by the user.

Locked Recovery Point

You can lock or unlock a recovery point as needed. After a recovery point is locked, data of the recovery point will not be automatically cleared or deleted.

Recovery Task

A recovery task helps you quickly restore data by specifying a CDP task and recovery point, and allows you to view the recovery progress and logs in a more friendly way.

Cryptography Security Compliance

The Cryptography Security Compliance service provides applications with cloud security capabilities based on commercial cryptography, meeting the requirements of commercial cryptography application security assessments.

HSM Pool

An HSM pool is a logical group of hardware security modules (HSMs) and is used to provide unified cryptography services such as signature validation and encryption.

HSM

A hardware security module (HSM) is a dedicated device that encrypts, decrypts, and authenticates information by using the cryptographic technology.

Platform Cryptography Security Compliance

Enables the Cloud to meet the requirements of Cryptography Security Compliance through the cryptography capabilities provided by HSM pools.

Certificate Login

Authenticates the identity of a user by using a UKey device.

Data Protection

Protects important data on the Cloud to ensure the data confidentiality and integrity.

Scheduled Job

A scheduled job defines that a specific action be implemented at a specified time based on a scheduler.

Scheduler

A scheduler is used to schedule jobs. It is suitable for business scenarios that last for a long time.

Tag

A tag is used to mark resources. You can use a tag to search for and aggregate resources.

Migration Service

The Cloud provides V2V migration service that allows you to migrate VM instances and data from other virtualized platform to the current cloud platform.

V2V Migration

V2V Migration allows you to migrate VM instances from the VMware or KVM platform to the current cloud platform.

V2V Conversion Host

A V2V conversion host is a host in the destination cluster that you need to specify during V2V migration to cache VM instances and data when you implement V2V migration. After the VM instances and data are cached in the V2Vconversion host, they are migrated to the destination primary storage.

User

A user is a natural person that constructs the most basic unit in Tenant Management.

User Group

A user group is a collection of natural persons or a collection of project members. You can use a user group to grant permissions.

Role

A role is a collection of permissions that can be granted to users. A user that assumes a role can call API operations based on the permissions specified by the role. Roles are categorized into platform roles and project roles.

3rd-Party Authentication

The 3rd-party authentication service provided by the Cloud. It supports seamless access to 3rd-party authentication systems. Through the service, related users can directly log in to the Cloud and manage cloud resources. Currently, AD/LDAP/OIDC/OAuth2/CAS servers can be added.

Project

A project is a task that needs to be accomplished by specific personnel at a specified time.
In Tenant Management, you can plan resources at the project granularity and allocate an independent resource pool to a project. The word **Tenant** in Tenant Management mainly refers to projects. A project is a tenant.

Project Member

A project member is a member in a project who is granted permissions on specific project resources and can use the resources to accomplish tasks. Project members include the project admin, project managers, and normal project members.

Process Management

Process management is part of ticket management that manages the processes related to the resources of projects. Processes can be categorized into default processes and custom processes .

My Approvals

In the Cloud, only the administrator and project administrators are granted approval permissions. the administrator and project administrators can approve or reject a ticket. If a ticket is approved, resources are automatically deployed and allocated to the specified project.

Bills

A bill is the expense of resources totaled at a specified time period. Billing is accurate to the second. Bills can be categorized into project bills, department bills, and account bills.

Pricing List

A pricing list is a list of unit prices of different resources. The unit price of a resource is set based on the specification and usage time of the resource.

Console Proxy

Console proxy allows you to log in to a VM instance 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 Cloud. AccessKey pairs shall be kept confidential.

IP Blocklist/Allowlist

An IP blocklist or allowlist identifies and filters IP addresses that access the Cloud. You can create an IP allowlist or blocklist to improve access control of the Cloud.

Application Center

Application Center allows you to add third-party applications to the Cloud and then access the applications by using the Cloud. It extends the functionality of the Cloud.

Sub-Account Management

A sub-account can be created by the admin or synced from a third-party authentication system and is managed by the admin. Resources created under a sub-account are managed by the sub-account.

Theme and Appearance

You can customize the theme and appearance of the Cloud.

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 cloud and use the collected logs to locate errors and exceptions. This makes your O&M more efficient.

Global Setting

Global Setting allows you to configure settings that take effect on the whole platform.

Scenario Template

Scenario Template provides multiple templates that encapsulate scenario-based global settings . You can apply a template globally with one click based on your business needs. This improves your O&M efficiency.

HA Policy

HA Policy is a mechanism that ensures sustained and stable running of the business if VM instances are unexpectedly or scheduled stopped or are errored because of errors occurring to compute, network, or storage resources associated with the VM instances. By enabling this feature, you can customize VM HA policies to ensure your business continuity and stability.

Time Management

Manages the Cloud system time and allows you to configure time servers for the Cloud. After you configure NTP time servers for the Cloud, the clock of the time servers is synced with all nodes of the Cloud.