Cloud–Virtual–Machine

Product Documentation

Document Contents

General

Last Updated At: 2025-08-04 16:48:35

Product Introduction
Features and Advantages
Features and Advantages

Last Updated At: 2025-08-12 21:08:34

Comprehensive CVM offers you a comprehensive range of services. ●Multiple Machine Models: -Standard Type (Suitable for small and medium-sized Web applications and small and medium-sized databases). -MEM-Optimized Type (Suitable for applications that require a lot of memory operations, searches, and calculations). -Big Data Type (Suitable for throughput-intensive applications such as Hadoop distributed computing, massive log processing, distributed file systems, and large data warehouses). -Heterogeneous Type (suitable for high-performance applications such as deep learning, scientific computing, video codec, graphics workstations, etc.). Flexible We are committed to building the most flexible CVM management platform in the industry, providing the following capabilities: ●Hardware Configuration: CVMs based on CBS allow instant upgrading/downgrading of hardware configurations. ●Disk Change: CVMs based on CBS allow instant scale-out of disks. ●Network Bandwidth: CVM allows instant upgrading/downgrading bandwidth. ●Operating System: The CVMs can switch between Windows and Linux systems at any time. ●Image Type: Public images (various Linux and Windows operating system types), custom images (images created by users through the image creation feature), and support for cross-regional adjustment and image replication. ●Customizable Network Architecture: VPC offers users independent network environments, tailored IP ranges and addresses, as well as personalized routing policies. Provides port-level access control to achieve comprehensive network logical isolation. Reliable Committed to building the most reliable CVM in the industry. ●CVM Reliability: Host service availability is 99.95%, data reliability is 99.9999999%. Supporting seamless failover migration, data snapshots, automatic alarms, and other features to safeguard your servers. ●CBS Policy: Provides a three-replica professional storage policy to eliminate single points of fault and ensure data reliability, allowing you to safely store your data in the cloud without worrying about data loss. ●Stable Network Architecture: Mature network virtualization technology and network interface binding technology ensure high network availability. Running in T3+ and above data centers ensures the reliability of the operating environment, freeing you from worries about network availability. Fast Whether in terms of user operation or CVM performance, we are committed to providing fast and convenient services. ●Easy and Fast Operation: You can easily obtain one, hundreds, or even thousands of server instances in just a few minutes. You can purchase, configure, scale, and manage your services with one click. ●Excellent Private Network: Featuring seamless interconnectivity among IDCs within the same region, with an underlying network capable of reaching speeds up to 10 Gbps or 1 Gbps, we guarantee the quality of private network communication. Security The cloud platform provides a variety of solutions to ensure the security of CVMs, and provides backup and rollback mechanisms to ensure data security. ●Multiple ways to remotely

log in to the CVM: Provides multiple log-in methods, including key log-in, password log-in, VNC log-in, etc. ●Rich Security Services: Provides Anti-DDoS, DNS hijacking detection, intrusion detection, vulnerability scanning, web Trojan detection, log-in protection, and other security services to protect your server. ●Free Cloud Monitor: Supports multiple real-time alarms. ●Recycle Bin Protection Mechanism: Supports cloud services to enter the recycling bin for a period of time before release, avoiding major impacts such as data loss caused by immediate termination. ●Customized Access Control: Customizes host and network access policies through security groups and network ACLs and flexibly sets different firewalls for different instances. The security service has the following characteristics: ●Comprehensive Security Guard Provides integrated security services for CVMs, including security checks (vulnerability scanning, Trojan horse detection, website backdoor detection, port security detection, etc.) and security defense (Anti-DDoS, intrusion detection, and access control to ensure data security and user privacy). ●Real-time Alarm and Regular Analysis 24/7 security service, discover vulnerabilities as soon as possible and notify you in real time for free. ●Free, Convenient and Secure No need to purchase expensive security equipment for your cloud service, you can enjoy Cloud Security service for free when you purchase cloud service. One-click activation, zero deployment, convenient and simple. ●Professional Team, Reliable Guarantee Cloud Security is created by a security team with many years of security experience and training. It provides professional security services to cloud service users and is worthy of your trust. Easy to Use Officially certified application software and Ops tools help you to operate and maintain conveniently, so you no longer have to worry about management tools. ●CVM provides a web-based user interface, the console, which can be used to start, adjust configurations, and reinstall the system on CVM instances just like physical machines. If you have registered a Cloud Platform account, you can log in to the CVM console directly to operate your CVM. ●CVM provides an API system. You can use the API to conveniently integrate the CVM with your internal monitoring and operation systems to achieve a fully automated business Ops system that is close to business needs. These requests are HTTP or HTTPS requests. For more information about CVM API operations, see the API documentation.

Snapshot

Snapshot

Last Updated At: 2025-08-12 21:08:34

Snapshot is an efficient data backup mechanism, particularly suitable for data protection in cloud computing environments. It captures the data state at a specific point in time, creating a fully usable copy that is independent of the original data's lifecycle. Even if the original data becomes damaged or lost, it can be recovered from the snapshot. Snapshots employ incremental backup strategies, recording only the data parts that have changed since the last snapshot. This approach not only reduces storage resource usage but also speeds up the creation of snapshots. Additionally, new cloud disks can be quickly deployed based on snapshots, inheriting the same data content for rapid test environment setup or disaster recovery scenarios. It's worth noting that snapshots typically have regional restrictions, meaning they can only be used within the same geographic region where the original disk resides. **Description** ●Real-Time Replica of Online Data A snapshot is a fully usable copy of a CBS. When a problem occurs in a CBS for which a snapshot has been created, the snapshot can be used to quickly restore it to its original status. It is recommended that you create a snapshot of the relevant CBS before making major business changes so that data can be quickly restored if the business change fails. ●Persistent Backup of Key Milestones Snapshots can serve as persistent backups of business data, preserving milestone status of business data. ●Rapid Business Deployment You can use business snapshot files to quickly clone multiple CBSs to quickly deploy servers. **Use Cases** Snapshot is a convenient and efficient data protection service and is recommended for the earlier business scenarios: ●Daily Data Backup You can use snapshots to regularly back up important business data to mitigate the risk of data loss caused by misoperation, attacks, or viruses. ●Quick Data Recovery You can create one or more snapshots before performing major operations such as changing the operating system, upgrading application software, or migrating business data. If any problems occur during the change operation, the business data can be restored in time through the created snapshot. ●Multi-replica Application of Production Data You can create production data snapshots to provide near-real-time production data for applications such as data mining, report query, development and testing. ●Rapid Deployment Environment You can create a snapshot of a CVM and use the snapshot to create a custom image. You can create one or more instances from the created image to quickly deploy CVMs in the same environment in batches, saving time on repeated configuration. **Snapshot Type** ●Manual Snapshot Manually create a snapshot of the CBS data at a certain point in time. This snapshot can be used to quickly create more CBSs with the same data, or to restore the CBS to the status at that point in the future. For detailed operations, see Creating a Snapshot. ●Periodic Snapshot As your business continues to evolve, periodic snapshots can be used to provide continuous backup capabilities. You only need to formulate a backup policy and associate it with the CBS to achieve continuous backup of the CBS data within a certain period, greatly improving data security.

Instance Overview

Last Updated At: 2025-08-12 21:08:34

Introduction to the Instance Instance, which can be regarded as Cloud Virtual Machine (CVM), contains the most fundamental computing components such as CPU, memory, operating system, network, and disk. CVM instances can offer secure, reliable and elastic computing services in the cloud to fulfill computing requirements. The instances can expand or contract computing resources in real time as the business needs change. They can significantly reduce the company's software and hardware procurement costs and simplify IT operation and maintenance work. Different instance types offer different computing and storage capabilities and are applicable to different use cases. Users can select the computing capability, storage and network access mode of the instance according to the scale of the services they need to offer. For more instance types and applicable scenes, please see CVM Instance Configuration. After the instance is started, the user can use it like a traditional computer, and the user has full control over the instance. Image of Instance Image It is a template for configuring the Cloud Virtual Machine software (operating system, pre-installed programs, and etc.). An image provides all the information required to start a Cloud Virtual Machine instance. Requires users to start instances through images. The image can start multiple instances for users to use repeatedly. In layman's terms, an image is the installation disk of Cloud Virtual Machine. The images include the following types. ●Public images: Available to all users, and applicable to most mainstream operating systems. ●Custom images: Only available to the creator and users that share the images, and the images are created by existing running instances or imported from external sources. ●Shared images: Images offered by other users, which can only be used to create instances. For more information about images, see Image Overview and Image Type. Storage of Instance The storage of the instance is similar to that of an ordinary Cloud Virtual Machine, and it includes a system disk and a data disk. ●System disk: Similar to the C Drive in the Windows system. The system disk contains a complete copy of the image used for starting the instance, as well as the instance operating environment. Before starting up, you must select a system disk size that is larger than that of the used image. ●Data disk: Similar to the D Drive and E Drive under the Windows system. The data disk stores the user data and supports free expansion, mounting and unmounting. Both system disks and data disks can use different storage types provided by the cloud platform。 For more information, see Storage Overview. Instance Security The instance security protection means provided by Cloud Platform are as follows. ●Policy control: When the same set of cloud resources needs to be controlled by multiple different accounts, users can use policy control to manage access permissions to cloud resources. ●Security group: Control access by using security groups to allow trusted addresses to access instances. ●Log in control: Use SSH keys whenever possible to log in to the user's Linux instance, and use password to log in to the instance that needs to change the password from time to time.

Specification

Last Updated At: 2025-08-12 21:08:34

When creating a Cloud Virtual Machine on the cloud platform, the instance type specified by the user determines the host hardware configuration of the instance. Each instance type provides different computing, memory and storage functions. Users can select an appropriate instance type based on the scale of applications they need to deploy. These instance families consist of different combinations of CPU, memory, storage, heterogeneous hardware, and network bandwidth, giving you the flexibility to choose the appropriate resources for your application.

Lifecycle

Last Updated At: 2025-08-12 21:08:34

The lifecycle of a Cloud Virtual Machine instance refers to all states from the start of the instance to its release. By managing instances properly from enabling to disabling instances, applications running on instances can provide services efficiently and economically. Instance State ●The states of an instance are as follows.

Status Name	Status Attributes	Status Description
Creating	Intermediate State	After the instance is created, it enters the status before it is running.
Running	Steady State	The instance is running properly. Your business can be run on this instance.
Restarting	Intermediate State	The instance enters the state before it is running after being restarted by the console or through the API. If this state is maintained for a long time, exceptions may occur.
Resetting	Intermediate State	The instance enters the state before it is running after the reinstall operation or the reset disk operation is performed in the console or through the API.
Shutting down	Intermediate State	The instance enters the state before it is shut down after being shut down by the console or through the API. If this state is maintained for a long time, it indicates an exception. Forced shutdown is not recommended.
Shut Down	Steady State	The instance is stopped normally. An instance in a shutdown state cannot provide services. Some attributes of the instance can only be modified when the system is shut down.
Terminating	Intermediate State	The instance has not been terminated 7 days after its expiration or after the user actively performs the termination.
Released	Steady State	The release operation is completed. The instance does not exist, and services are unavailable, and the data is completely cleared.

Launch an Instance ●The instance enters the creating state after the it is started. The hardware for the instance being created is configured according to the specified instance specifications, and the system starts the instance by using the image specified at the startup. ●The instance enters the running state

after it is created. Instances in the running state enable normal connections and access services. Restart an Instance It is recommended to choose to restart an instance by using Console or API instead of running the restart command of the operating system in the instance. ●The instance enters the restarting state after the restart an instance operation is performed. ●Restarting an instance is equivalent to restarting a computer. After the instance is restarted, it still retains its public IP address, private IP address, and all data on its hard disk. ●Restarting an instance usually takes from tens of seconds to several minutes, depending on the instance configuration. For more information about restarting an instance, see Restart an Instance. Shutdown an Instance Users can use the console or API to shut an instance down. ●Shutting down an instance is equivalent to turning off your computer. ●After the instance is shut down, it will no longer provide services, but the billing will not stop. ●An instance being shut down is still displayed in the console. ●Shutting down an instance is a prerequisite for some configuration operations, such as adjusting hardware configuration and resetting passwords. ●The shutdown operation itself does not change the public network IP address, private IP address of the Cloud Virtual Machine and all the data on its hard disk. Terminate or Release an Instance When users no longer need a Cloud Virtual Machine instance, they can terminate and release the instance. The termination and release can be achieved by using the Console or API.

Storage

Storage Overview

Last Updated At: 2025-08-13 14:18:05

We provide various types of data storage devices that are flexible, economical, and easy to use for Cloud Virtual Machine instances. Different storage devices have different performance and prices and are suitable for different use cases. Storage device classification Storage devices can be divided into the following categories according to different classification dimensions.

Division dimensions	Classify	Note
Storage Media	SSD Hard Disks	The storage medium is a solid state drive (SSD). It is characterized by excellent performance in IOPS and reading and writing speed, and can achieve up to 20 times the IOPS and 16 times the throughput of ordinary hard disks. The price is higher than ordinary hard disks.
Application Scenario	System Disk	A collection of systems used to store, control, and schedule Cloud Virtual Machine operations by using images.
	Data Disk	Used to store all user data.
Architecture Patterns	Cloud Block Storage	CBS is a flexible, highly available, reliable, low-cost, and customizable network block device that can serve as an independently scalable hard drive for CVM. It provides block-level data storage with a three-copy distributed mechanism to ensure data reliability for CVM. Choosing CBS for CVM enables hardware, disk, and network adjustments.
	Local Disk	The local disk is a part of the local storage of the physical machine where the CVM instance is located. It is a storage area of the local storage of the physical machine where the CVM instance is located. Data access can achieve lower latency, but there is a risk of single point failure of data. The hardware (CPU, memory, and disks) upgrade is not available for the CVM that uses local disks, but the bandwidth upgrade is supported.
	COS	Cloud Object Storage is a data storage device located on the Internet that supports retrieval of data from a Cloud Virtual Machine instance or anywhere on the Internet, thereby reducing storage costs. Not suitable as a storage medium for low-latency, high-IO scenes.

Block storage device mapping Each instance has a system disk to ensure basic operating data, and more data disks can be mounted to the instance. Instances use block storage device mapping to map these storage devices to locations that they can recognize. Block Storage is a storage device that is divided into blocks of bytes and supports random access. The cloud platform supports two types of block storage devices, that is, local disks and Cloud Block Storage.

Cloud Block Storage

Last Updated At: 2025-08-12 21:08:34

Cloud Block Storage (CBS) provides you with persistent data block-level storage services for Cloud Virtual Machines. ●The data in the Cloud Block Storage is automatically stored in multiple copies within the availability zone, avoiding the risk of single point failure of data and providing up to 99.9999999% data reliability. ●Cloud Block Storage provides disk instances of various types and specifications to meet the requirements of stable and low-latency storage performance. ●Cloud Block Storage can be mounted/unmounted on instances in the same availability zone, and storage capacity can be adjusted within minutes to meet flexible data needs. You only need to pay a fair price for the configured resources to enjoy the above features. Typical use cases ●If the Cloud Virtual Machine is found to have insufficient hard disk storage capacity during use, you can purchase one or more Cloud Block Storages and mount them to the Cloud Virtual Machine to meet the storage capacity requirements. ●No additional storage space is required when purchasing a Cloud Virtual Machine. When storage needs arise, the storage capacity of the Cloud Virtual Machine can be scaled up by purchasing a Cloud Block Storage. ●When there is a demand for data exchange among multiple Cloud Virtual Machines, the demand can be fulfilled by unmounting the Cloud Block Storage (data disk) and remounting it to other Cloud Virtual Machines. ●You can purchase multiple Cloud Block Storages and configure LVM (Logical Volume Manager) logical volumes to break the storage capacity upper-limit of a single Cloud Block Storage. ●You can purchase multiple Cloud Block Storages and configure RAID (Redundant Array of Independent Disks) policies to break the I/O capacity upper-limit of a single Cloud Block Storage. Lifecycle The lifecycle of non-elastic cloud disks completely follows that of the cloud server; they are applied for together with the cloud server and used as system disks, and do not support mounting and unmounting operations. The lifecycle of elastic cloud disk is independent of the Cloud Virtual Machine instance and is not affected by the instance operation. You can attach multiple Cloud Block Storages to a single instance, or detach a Cloud Block Storage from an instance and attach it to another instance.

COS

Last Updated At: 2025-08-12 21:08:34

Cloud Object Storage (COS) is a distributed storage service provided by cloud platforms to store massive amounts of files, enabling users to store and access data anytime over the network. CVM users can store and retrieve data via the instance or anywhere on the Internet. COS redundantly stores user data across multiple regions and allows multiple different clients or application threads to read or write this data simultaneously. COS provides CVM users with a highly scalable, low-cost, reliable, and secure data storage solution.

Local Disk

Last Updated At: 2025-08-13 14:18:06

Local Disk Overview A local disk is a storage device that is located on the same physical server as a Cloud Virtual Machine (CVM) instance and has the characteristics of high reading and writing speed IO and low latency. The local disk is a part of the local storage of the physical machine where the CVM instance is located. It is a storage area of the local storage of the physical machine where the CVM instance is located. The local disk is a part of a single physical machine. The data reliability depends on the reliability of the physical machine, and there is a risk of single point failure. Lifecycle: The lifecycle of a local disk is the same as the CVM instance to which it is mounted. In simple terms, the local disk starts and ends along with the startup and shutdown of the CVM it is mounted to. Application: Local disks can only be initiated alongside the startup of the cloud server. Therefore, applying for local disks can only be done when applying for cloud server instances. For more information on applying for cloud servers, see Apply and Launch Instance. Note: Cloud servers equipped with local disks do not support hardware upgrades (such as CPU and memory), only bandwidth upgrades are supported. Type Local disks originate from the local storage of the physical machine hosting the cloud server. Depending on the medium, they can be categorized into Ordinary Local Disks and SSD Local Disks. Ordinary Local Disk

Specification	Application Strategy	Performance
System Disk	Fixed at 50GB, cannot be modified	Peak throughput above 40~100 MB/s, IOPS ranging from several hundred to 1000
Data Disk	Supports ordinary local disk specifications from minimum 10GB to maximum 1600GB (in increments of 10GB), and different hardware configurations allow selection of ordinary local disk specification limits that vary.	

SSD Local Disk The SSD local disk originates from the local storage of the physical machine hosting the cloud server. Such storage provides block-level data access capabilities using full SSD media, featuring low latency, high random IOPS, and high throughput I/O capabilities.

Specification	Application Strategy	Performance
System Disk	From 50G to 500G (cannot be expanded after initialization),	Peak throughput 250 MB/s Maximum random write IOPS up to 10000 (for 4K random write depth 32) Maximum random read IOPS up to 75000 (for 4K random read depth 32) Access latency less than 3 ms
Data Disk	Supports SSD local disk specifications from minimum 10GB to maximum 16000GB (in increments of 10GB), and different hardware configurations allow	

	selection of ordinary local disk specification limits that vary.	

SSD local disks are suitable for use in the following scenarios: Low Latency: Providing microsecond–level access delay. Distributed Applications: I/O–intensive applications such as NoSQL databases, MPP data warehouses, and distributed file systems, which inherently possess distributed data redundancy capabilities. Large–scale Online Application Logs: Large–scale online applications generate massive volumes of log data, necessitating high–performance storage, while log data has relatively lower requirements for storage reliability. Single–point Risk: There is a risk of single–point failures, recommending implementing data redundancy at the application level to ensure data availability.

Image Overview

Last Updated At: 2025-08-13 14:29:45

What Is Image? Image is a template for CVM software configuration (operating system, pre-installed programs, etc.), which provides all the information required to start a CVM instance. An image provides all the information required to start a Cloud Virtual Machine instance. Requires users to start instances through images. The image can start multiple instances for users to use repeatedly. In layman's terms, an image is the installation disk of Cloud Virtual Machine. Image Type The images include the following types. ●Public Image: Available to all users, covering most mainstream operating systems. ●Custom Image: Only available to creators and shared objects, created from an existing running instance or imported from outside. ●Shared Image: Images shared by other users can only be used to create instances. Image Deployment Vs Manual Deployment

	Image Deployment	Manual Deployment
Deployment Time	3 to 5 minutes	1 day – 2 days
Deployment Process	Quickly create a suitable CVM based on the used solutions.	Select the appropriate operating system, database, application software, plugins, etc., and install and debug them.
Security	Except for shared images, which require users to identify their sources, other public images and custom images have been tested and reviewed.	Depends on the level of development and deployment personnel.
Usage	Public Image: genuine operating system, including the initialization components provided by Cloud Platform.Custom Image: Quickly create the same software environment as an existing CVM or back up the environment.Shared Images: Quickly create the same software environment as other users' existing CVMs.	Completely self-configured, with no basic settings required.

Image Application ●Deployment of Specific Software Environment Using shared images or custom images can help quickly set up specific software environments, eliminating the tedious and time–consuming tasks of self–configuring environments and installing software. They can also meet various personalized needs such as website building, application development, and visualization management, making CVMs ”ready to use” immediately, saving time, and providing convenience. ●Batch Deployment of Software Environment By creating an image for a CVM instance with a deployed environment, and then using the image as the operating system when creating CVM instances in batches, the CVM instance will have the same software environment as the previous CVM instances after it is successfully created, thereby achieving the purpose of batch deployment of software environments. ●Server Operating Environment Backup Create an image backup of the operating environment for a CVM instance. If the CVM instance cannot run normally due to software environment damage during use, you can use image recovery. Image Lifecycle The following diagram summarizes the lifecycle of a custom image. After creating or importing a new custom image, users can use it to launch new instances (users can also launch instances from existing public images). Custom images can be copied to other regions of the same account and become independent images in that region. Users can also share custom images with other users.

Type

Last Updated At: 2025-08-14 09:28:32

Public Image Public Image is an image officially provided, supported, and maintained by Cloud Platform, including the basic operating system and initialization components provided by Cloud Platform, and can be used by all users. Public Image Features: ●Operating System Types: free choice (e.g.: based on Linux type system or Windows type system), and regularly updated. ●Software Support: Integrate the software packages provided by Cloud Platform (such as API, etc.), and support multiple versions of common software such as Java, MySQL, SQL Server, Python, Ruby, and Tomcat, and their full permissions. ●Security: The operating systems provided are completely legal and compliant, and all use official genuine operating systems. Produced by the professional security Ops team, it has undergone rigorous testing and offers optional built-in security components. ●Restrictions: no use limits.

Custom image Custom Image It is an image created by the user through the image creation feature or imported through the image import function. Only the creator and the sharer can use it. Custom Image Features: ●Use Cases: Create an image for a CVM instance that has an application deployed, so that you can quickly create more instances with the same configuration. ●Feature Support: Supports users to create, copy, share, and terminate. ●Restrictions: Each region supports up to 10 custom images. Shared Image Shared Image refers to a custom image that is shared by another user to the current user through the image-sharing feature. The shared image will be displayed in the same region as the original image of the shared user. Shared Image Features: ●Use Cases: Help other users quickly create CVMs. ●Feature Support: Shared images can only be used to create CVMs. Other operations such as modifying the name, copying, and sharing are not allowed. ●Security: Shared images are not reviewed by Cloud Platform and may pose security risks. Therefore, it is strongly recommended not to accept images from unknown sources. ●Restrictions: Each custom image can be shared with up to 50 users. Image sharing only supports sharing to accounts in the same region.

Network and Security

Network and Security Overview

Last Updated At: 2025-08-13 14:40:28

The cloud platform provides network and security features to ensure that your instances can provide services internally and externally safely, efficiently, and freely. Encrypted Login Method Provides two encrypted login methods: password login and SSH key pair login. Users can freely choose between two methods for safely connecting to CVMs. Instances with Windows system do not support SSH key login. Network Access Cloud services on the same cloud platform can be accessed via the Internet or the private network. ●Internet Access: Internet access is a service provided by Cloud Platform to instances for public data transmission. Instances are assigned public IP addresses to enable communication with other computers on the network. ●Internal Network Access: Internal network access refers to local area network (LAN) service, which is a service provided by Cloud Platform to instances through private IP addresses, enabling completely free private network communication within the same region. Network Environment network environment can be divided into: basic network and Virtual Private Cloud (VPC). ●Basic Network: The basic network is a public network resource pool for all users on Cloud Platform. Suitable for users who are just beginning to understand and use Cloud Platform. ●VPC: A VPC is a logically isolated network space that you customize on Cloud Platform. Instances in a VPC can be launched in a preset, custom IP range and isolated from other users. Suitable for users who are familiar with network management system. Security Group A security group is a stateful virtual firewall with a packet filtering function. It is used to set network access control for one or more CVMs and is an important network security isolation method provided by Cloud Platform. You can control access permission to your instances using the earlier methods: ●Create multiple security groups and assign different rules to each security group. ●Each instance is assigned one or more security groups. Cloud Platform will use these rules to determine which traffic can access the instance and which resources the instance can access. ●Configure the security group so that only specific IP addresses or specific security groups can access the instance. Elastic IP (Currently Unsupported) Elastic IP (EIP), also known as elastic IP address. It is a static IP address designed for dynamic cloud computing. In the earlier scenarios, it is recommended to use EIP: ●The instance may go down due to uncontrollable reasons, and an alternative instance with the same IP address is required to ensure access. ●The instance does not have a public IP address and requires a static IP address. Elastic Network Interface Elastic Network Interface (ENI) is an elastic network interface bound to CVMs in a VPC, which can be freely migrated among multiple CVMs. ENI is very helpful in configuring and managing networks and building highly reliable network solutions.

Private Network Service

Last Updated At: 2025-08-14 09:28:16

Private network services are local area network (LAN) services, and cloud services access each other via internal links. Cloud services can be accessed via the Internet, and can also access each other through the private network on a cloud platform. The Cloud Platform IDCs are interconnected with underlying 10Gbps/1Gbps networks, providing high-bandwidth and low-delay private network communication services. Additionally, private network communication within the same region is completely free, enabling you to flexibly build your network architecture. Private IP Address Overview The private IP address is an IP address that cannot be accessed through the Internet and is the implementation form of Cloud Platform private network service. Each instance has a default network API (i.e., eth0) with a private IP address that can be automatically allocated by Cloud Platform or defined by the user (only in a VPC environment). Attribute ●The private network service has user attributes, and different users are isolated from each other, which means that by default, one user cannot access another user's cloud services via the private network. ●The private network service has regional attributes, and different regions are isolated from each other, which means that by default, one cannot access cloud services in different regions under the same account via the private network. Applicable Scenario The private IP address can be used for Cloud Load Balancer, private network access between CVM instances, and private network access between CVM instances and other cloud services (such as T DB). Address Allocation Each CVM instance is assigned a default private IP address when it is started. For different network environments, the private IP address is also different: ●Basic network: The private IP address is automatically assigned by Cloud Platform and cannot be changed. ●VPC: Currently, Cloud Platform VPC CIDR supports the use of any of the three major private IP ranges. The IP address range is as follows, and the mask range must be between 16 to 28: -10.0.0.0 - 10.255.255.255 -172.16.0.0 - 172.31.255.255 -192.168.0.0 - 192.168.255.255 Private Network DNS DNS Server Address The private network DNS service is responsible for domain resolution. If the DNS configuration is incorrect, the domain will be inaccessible. The cloud platform provides reliable private network DNS servers in different regions. Private Network DNS Settings When network resolution errors occur, users can manually set up intranet DNS settings. The setup method is as follows:

Obtain the instance's Private Network IP address using the console 1.Log in to the cloud server console. 2.In the list of cloud servers under your account, move the mouse over the internal IP of the cloud server; when the copy button appears, click it to copy the internal IP.

Elastic Network Interface

Last Updated At: 2025-08-13 15:07:19

Elastic Network Interface (ENI) is an elastic network interface bound to CVMs in a VPC, which can be freely migrated among multiple CVMs. ENI is very helpful in configuring and managing networks and building highly reliable network solutions. An ENI combines VPC, availability zone, and subnet attributes and can only be bound to CVMs in the same availability zone. A CVM can bind multiple ENIs, and the specific number of bindings will depend on the host specifications. Relevant Concepts ●Primary NIC and Secondary NIC: The NIC that is created simultaneously with the creation of a CVN in a VPC is the primary NIC, while those NICs created by users are considered secondary NICs. The primary NIC does not support binding or unbinding, while the secondary NICs do support binding and unbinding. ●Primary Private IP Address: The primary private IP address of the ENI is randomly assigned by the system or set by the user when the ENI is created. The primary private IP address of the primary NIC can be modified, but the primary private IP address of the secondary NIC cannot be modified. ●Secondary Private IP Address: The secondary private IP addresses bound to an ENI except the primary IP can be configured by users when creating or editing an ENI, and support binding and unbinding. ●EIP: Bind one by one with the private IP address on the ENI. ●Security Group: An ENI can be bound with one or multiple security groups. ●MAC Address: The ENI has a globally unique MAC address. Use Cases ●Isolation of Private Network, Public Network, and Management Network: The network deployment of important businesses generally requires the isolation of the private network, public network, and management network for data transmission, and ensures data security and network isolation between networks through different routing policies and security group policies. Just like a physical server, you can bind three ENIs in different subnets to the CVM to achieve triple network isolation. ●High-Reliability Application Deployment: The key components in the system architecture need to be deployed in a multi-machine hot standby mode to ensure the high availability of the system. Cloud Platform provides ENIs and private IP addresses that can be flexibly bound and unbound. You can configure the disaster recovery settings of Keepalived to achieve high-availability deployment of key components. Use Limits Depending on CPU and memory configurations, the number of elastic NICs that can be bound to a cloud server and the number of internal IPs per NIC vary significantly. The quota for NICs and single-NIC IP numbers is shown in the following table:

Cloud Server Configuration	Elastic NIC Number	NIC-bound IP Number
CPU: 1 core Memory: 1 GB	2	2
CPU: 1 core Memory: > 1 GB	2	6
CPU: 2 cores	2	10
CPU: 4 cores Memory: < 16 GB	4	10

Cloud Server Configuration	Elastic NIC Number	NIC-bound IP Number
CPU: 4 cores Memory: > 16 GB	4	20
CPU: 8-12 cores	6	20
CPU: > 12 cores	8	30

Configuration Operation Guide If the cloud server needs to use an elastic network interface card (ENI), please follow the configuration steps below to complete the relevant tasks:1. Create an Elastic Network Interface Card.2.Bind the Elastic Network Interface Card to the Cloud Server.3. Configure the cloud server and private network routing table, see Private Network Routing and Security Configuration for details.4. Perform the configuration of the ENI within the cloud server.5. Allocate an Internal IP Address. Depending on your requirements, refer to Allocate Internal IP Address (within the Cloud Server System) or Allocate Internal IP Address (Qcloud Console). For more operations related to elastic network interface cards, please refer to the Elastic Network Interface Card Operation Guide. API Overview The following table shows the APIs related to ENI and CVM. For more ENI-related operations, see ENI API Overview.

API Feature	Action ID	Feature Description
Create an ENI	CreateNetworkInterface	Create an ENI
Apply for a private IP address using an ENI	AssignPrivateIpAddresses	Apply for a private IP address using an ENI
Bind an ENI to a CVM	AttachNetworkInterface	Bind an ENI to a CVM

Login Password

Last Updated At: 2025-08-13 15:07:19

To ensure the security and reliability of instances, the cloud platform offers two encrypted login methods: password login and SSH key pair login. Users of cloud servers running different operating systems can refer to the Customized Configuration section of both Windows Cloud Servers and Linux Cloud Servers respectively, where they can choose their preferred encryption method. A password serves as the exclusive login credential for each cloud server instance. Anyone who possesses the instance login password can remotely access the cloud server instance via the public IP address allowed by the security group. Therefore, we recommend that you use a relatively secure password, keep it safe, and change it periodically. Setting Initial Password • For users selecting ”Automatically Generate Password,” the initial password will be sent via the in-console message center. • For users choosing ”Set Password,” the custom password becomes the initial password. 1.When customizing the configuration of a cloud server, users can select the login method during the Hostname and Login Method setup phase, with the default being ”Set Password.” 2.Following the specified password character restrictions, enter the host password and confirm it before clicking Apply Now. Once this is done, the initial password setting is successful, pending allocation of the cloud server instance. Password Character Restrictions: o For Linux cloud servers, passwords must be between 8 and 16 characters long, including at least two categories among lowercase letters (a-z), uppercase letters (A-Z), digits (0-9), and special characters o For Windows cloud servers, passwords must be between 12 and 16 characters long, including at least three categories among lowercase letters (a-z), uppercase letters (A-Z), digits (0-9), and special characters Viewing Password Passwords automatically generated will be sent to the in-console message center. Click on the relevant message to view the initial password. Access the in-console messages by clicking on the envelope icon located at the top right corner of the Cloud Server Console. Resetting Password Note: Password resetting is only possible when the cloud server is powered off. If the cloud server is running during password reset, it will force shutdown, potentially leading to data loss or damage to the file system. 1.Log in to the Cloud Server Console. 2.Power off the cloud server that requires password reset. 3.Open the password reset dialog box. o For individual instances that are powered off, click ”More” -> ”Reset Password” in the right-side operation panel. o For multiple instances, check all hosts requiring password reset, then click ”Reset Password” at the top of the list to batch update the host login passwords. Reasons for instances unable to reset passwords will be displayed. 4.In the pop-up window for password reset, enter the new password and confirm it before clicking Next. 5.Wait for the reset to succeed. You will receive a success notification via the in-console message center, allowing you to power on and use the cloud server with the new password.

SSH Key

Last Updated At: 2025-08-13 15:07:19

To ensure the security and reliability of the instance, two encrypted login methods are provided: password login and SSH key pair login. This document describes the configuration of SSH key pair login. When custom-configuring a Linux CVM, you can select the SSH key as the CVM encrypted login method. SSH Keys Overview The Cloud Platform recommends that you log in to Linux instances using SSH key pairs. An SSH key pair is a pair of keys generated through encryption algorithms. The SSH key pairs created by the cloud platform use RSA 2048-bit encryption, namely a public key and a private key: Public Key: After the SSH key pair is successfully generated, the cloud platform only stores the public key information. Private Key: You need to download and keep the private key properly. The private key can only be downloaded once, and the cloud platform will not save your private key. Anyone with your private key can gain access to your login information, so you need to keep your private key in a safe place. You can use a key pair to securely connect to the CVM. Using a key pair to log in to the CVM is more secure than using a regular password. You only need to specify a key pair when creating an instance, or bind a key pair after the instance is created, and then you can use the private key to log in to the Linux instance without entering a password. Features and Advantages Compared with the traditional username and password authentication method, using SSH keys has the following advantages: ●SSH key login authentication is more secure and reliable, and can eliminate the threat of brute force cracking. ●The SSH key login method is simpler. You only need to make simple configurations in the console and local client to remotely log in to the instance. You do not need to enter a password when you log in again. Use Limits ●Only Linux instances are supported. ●The cloud platform will not keep your private key information. Users need to click Download to obtain the private key within 10 minutes after creating the SSH key, and keep it properly. ●For data security considerations, loading the key needs to be done while the system is powered off. ●To improve the security of the cloud server, password login will be disabled by default after the instance is bound to a key pair. If you wish to use password login concurrently, please proceed to the cloud server console to reset the instance password.

Security Group

Last Updated At: 2025-08-14 09:27:50

Overview of Security Groups A security group is a stateful packet filtering virtual firewall used for setting network access control for single or multiple cloud servers. It's an important means of network security isolation provided by the cloud platform. • A security group is a logical grouping. You can add basic network cloud servers or elastic network card instances with the same network security isolation requirements within the same region into the same security group. • You can apply security group policies to filter inbound and outbound traffic for instances, which could be basic network cloud servers or elastic network card instances. • You can modify the rules of a security group anytime, and the new rules will take effect immediately. Security Group Templates Security groups support custom creation and template-based creation. Currently, three templates are provided: • Linux Open Port 22: Only exposes TCP port 22 for SSH login to the public internet; all internal network ports are open. • Windows Open Port 3389: Only exposes TCP port 3389 for RDP login to the public internet; all internal network ports are open. • All Ports Open: Exposes all ports to both the public internet and the internal network, which poses certain security risks. Security Group Rules Security group rules control the inbound traffic allowed to reach the instances associated with the security group, as well as the outbound traffic allowed to leave these instances (rules are processed top-down). By default, newly created security groups will drop (reject) all traffic. Binding a server to a rule-less security group blocks all traffic. For each rule in a security group, you can specify the following items: • Type: Choose system-defined templates or create custom rules. • Source or Destination: The source (for inbound rules) or destination (for outbound rules) of the traffic. Specify one of the following options: o Using CIDR notation, specify individual IP addresses. o Using CIDR notation, specify ranges of IP addresses (for example: 203.0.113.0/24). o Referencing a security group ID, choose one of the following security group IDs: □Current security group (indicates whether CVMs associated with this security group can/must communicate with each other) □Another security group in the same region. o Reference IP address objects or IP address groups defined in parameter templates. • Protocol and Port: Fill in protocol types and port ranges. You may also reference protocol-port pairs or protocol-port groups defined in parameter templates. • Policy: Allow or Deny. Note: o Referencing a security group ID is considered an advanced feature. The rules of the referenced security group won't be added to the current security group. o When configuring security group rules, if specifying a security group ID in the source/destination field, it refers exclusively to the private IP addresses bound to CVMs or elastic network cards associated with this security group ID, excluding public IP addresses. Security Group Priority • When multiple security groups are bound to an instance, priority is determined by lower numerical values indicating higher priority. • Within a security group, rule priority is determined by their position, where rules higher up have greater priority. Security Group Limitations • Security groups are distinguished by region and project; CVMs can only be bound to security groups in the same region and project. • Security groups apply to any CVM instances under network environments. • Each user can set

up to 50 security groups per region per project. • For inbound or outbound directions of a security group, access policies can each have up to 100 rules defined. • A CVM can join multiple security groups, and one security group can associate with multiple CVMs without limit on quantity. • Security groups attached to basic network cloud servers cannot filter packets coming from (or going to) CDBs or elastic caches (Redis and Memcached) on the cloud platform. If you need to filter traffic for such instances, you may use iptables to achieve this.

Feature Description	Quantity
Security Groups	50 / region
Access Policies	100 / inbound, 100 / outbound
Number of instances associated with a security group	Unlimited
Number of instances within a security group	Unlimited

Note: If you have a large number of instances that need mutual access, you can allocate them into multiple security groups and grant each other access through rule configurations referencing security group IDs, allowing mutual visits. Operation Guide You can manage security groups and their rules via the Cloud Server Console, including creating, viewing, updating, and deleting operations. Getting Started A security group is an instance–level firewall provided by the cloud platform, enabling control over inbound/outbound traffic for any cloud server. 1.Log in to the Cloud Server Console and click [Security Groups] in the left navigation pane. 2.Click the [Create] button, enter the name of the security group (e.g., my–security–group), choose template–based creation or custom creation, confirm the ingress/egress rules, then click [OK]. 3.On the right side of the security group list, click the [Add Instances] button, select the required associated virtual hosts, completing the operation of associating the security group with the virtual host. Alternatively, 4.You can also go to the list page of virtual hosts, view or modify the security groups already bound to a specific virtual host. Select the virtual host whose security group needs adjustment in the [Virtual Host] list page, then click [More] – [Configure Security Group] on the right, selecting the security group binding. Creating a Security Group 1.Access the console –> Security Groups. 2.In the left navigation pane, click [Security Groups]. 3.Click the [New] button. 4.Enter the name of the security group (for example: my–security–group) and provide a description. 5.Click [OK] to complete the creation process. Adding Rules to a Security Group 1.Access the console –> Security Groups. 2.In the left navigation pane, click [Security Groups]. 3.Select the security group to update, and click its [Security Group ID]. This displays detailed information about the security group in the details pane, along with tabs for inbound and outbound rules available for your use. 4.On the inbound/outbound rules tab, click [Edit]. Select the options for inbound/outbound rules from the tab, fill in the required information, and after completion, click [Save]. Configuring Security Groups for CVM Instances 1.Access the console –> Virtual Hosts. 2.In the left navigation pane, click [Virtual Hosts]. 3.In the operations bar on the right side of the

instance requiring security group configuration, click [More], then [Configure Security Group]. 4.In the Configure Security Group dialog box, select one or more security groups from the list, then click [OK]. Or 1.Access the console –> Security Groups. 2.Click [Security Groups] in the left navigation pane. 3.Select the security group to associate, and click the [Add Instances] or [Remove Instances] button in the operations bar. 4.In the pop–up window for adding/removing virtual hosts, add or remove the virtual hosts that need to be associated with this security group, then click [OK]. Importing and Exporting Security Group Rules 1.Access the console –> Security Groups. 2.In the left navigation pane, click [Security Groups]. 3.Select the security group to update, and click its [Security Group ID]. This displays detailed information about the security group in the details pane, along with tabs for inbound and outbound rules available for your use. 4.Select the options for inbound/outbound rules from the tab, then click the [Import Rules] button. If you already have existing rules, we recommend exporting your current rules first since importing new rules will overwrite the original ones; if you have empty rules originally, you can export a template first, edit the template file, and then import the file. Cloning a Security Group 1.Access the console –> Security Groups. 2.In the left navigation pane, click [Security Groups]. 3.Click the [Clone] button corresponding to the security group in the list. 4.In the Clone Security Group dialog box, select the target region and target project, then click [OK]. If the new security group needs to be associated with CVMs, reconfigure the security group accordingly. Deleting a Security Group 1.Access the console –> Security Groups. 2.In the left navigation pane, click [Security Groups]. 3.Click the [Delete] button corresponding to the security group in the list. 4.In the Delete Security Group dialog box, click [OK]. If the current security group is associated with CVMs, you must first dissociate the security group before deletion is possible. Differences between Security Groups and Network ACLs:

Security Group	Network ACL
Operates at the instance level (first line of defense)	Operates at the subnet level (second line of defense).
Supports both allow and deny rules	Supports both allow and deny rules
Stateful: Return traffic is automatically allowed without being affected by any rules.	Stateless: Return traffic must be explicitly allowed by the rules
Operations are applied to instances only when a security group is specified during instance launch or associated with an instance later on	Automatically applied to all CVM instances within the associated subnet

Security Group Cloud API The developer tool for security groups allows you to manage operations on security groups and configurations between security groups and CVM instances via cloud APIs.

Monitoring and Alarms

Monitoring and Alarms

Last Updated At: 2025-08-14 09:27:20

Monitoring and alarm is an important part to ensure high reliability, high availability and high performance of CVM. This document provides an overview of the monitoring and alarm features provided for the CVM. For more details, see Cloud Monitoring Product Documentation. Overview CVM monitoring and alarm is a management tool for real-time monitoring of CVMs. The monitoring and alarm feature can display the most comprehensive and detailed monitoring data, extracting key indicators from the CVM in real-time and presenting them in the form of monitoring charts. To provide you with a comprehensive understanding of your CVM's resource usage, performance, and health status. It also supports setting custom alarm thresholds and sending notifications based on your custom rules. Basic Features The console provides access to the following functionalities for CVM monitoring and alarm:

Module	Capability	Main Feature
Monitoring Overview	View the overall monitoring status of cloud services.	Provide an overview of the overall situation, alarm situation, and overall monitoring information.
Alarm Management	Support for user custom alarm thresholds	Currently supports CVM alarm settings
Cloud Services Monitoring	View detailed monitoring status of cloud services.	Current CVM monitoring view
Dashboard	View the monitoring view of cloud products	Current monitoring view of the cloud server
Custom Monitoring	View user-defined monitoring metric data.	Provide custom metric reporting and monitoring alarm services.
Traffic Monitoring	Monitor traffic usage.	View overall user bandwidth information.

Application Scenario Daily management scene: Log in to the cloud monitoring console to view the running status of each cloud monitoring. Timely handling of exception scene: When the monitoring data reaches the alarm threshold, an alarm message will be sent to notify you promptly, allowing you to acquire exception notification and check the cause of the exception. Timely scale-out scene: After the alarm rules for monitoring items such as bandwidth, connection counts, disk usage, etc. are set, you can conveniently

understand the current status of your cloud service and receive timely alarm notifications to scale out your service capacity as the business volume increases. Monitoring Content To monitor instance performance benchmarks, you should monitor at least the following monitoring items:

Monitoring Items	Monitoring Metrics	Explanation
CPU Utilization	cpu_usage	The CPU usage ratio is collected and reported by internal monitoring components on the server for more accurate data.
Memory Utilization	mem_usage	The ratio of actual memory used by the user to total available memory, excluding buffer and system cache occupied memory.
Private Network Outbound Bandwidth	lan_outtraffic	Average outgoing traffic per second on the intranet network card.
Private Network Inbound Bandwidth	lan_intraffic	Average incoming traffic per second on the intranet network card.
Outbound Bandwidth of Public Network	wan_outtraffic	Average outgoing traffic over the external network per second, calculated from the total traffic divided by 10 seconds, with the minimum granularity being 10 seconds.
Public Network Inbound Bandwidth	wan_intraffic	Average incoming traffic over the external network per second.
Disk Usage Rate	disk_usage	Disk utilization rate.
Disk I/O Waiting Time	disk_io_await	The average waiting time for each operation of hard disk I/O.

Monitoring Data □Monitoring interval: Currently, cloud monitoring provides multiple monitoring data statistical granularities of 1 minute, 5 minutes, 1 hour, and 1 day. The CVM can support 1 minute monitoring granularity, that is, data is counted every 1 minute, and the default interval is 5 minutes. □Data storage: Monitoring data with 1 minute, 5 minutes and 1 hour granularity is stored for 31 days, and monitoring data with 1 day granularity is stored for half a year. □Alarm display: The data is displayed in easy-to-read charts. The console integrates the monitoring data of all products, which is more conducive to users to obtain a holistic operation overview. □Alarm settings: You can set threshold values for

monitoring metrics, and when the conditions are met, alarm messages will be sent promptly to the concerned group. For details, see Creating an Alarm Policy. □Dashboard Settings: You can set up Dashboards for monitored metrics, which dynamically analyze the reasons for metric anomalies through the Dashboard. It also allows real-time observation of changes in metrics and timely resource expansion when necessary.

Access Control

Access Control

Last Updated At: 2025-08-13 15:58:47

Access Control Overview If you use CVM, VPC, databases and other services in the cloud platform, and these services are managed by different people but share your cloud account token, the following problems will occur: The risk of your key being compromised is high since multiple users are sharing it. The access permission of other users is not under control. They can introduce security risks caused by misoperations. In this case, you can use sub-accounts to allow different users to manage different services to prevent the problems above. By default, a sub-account has no permission to use CVM or CVM-related resources. Therefore, we need to create policies to allow sub-accounts to use the resources or permissions they need. Cloud access management (CAM) is a set of web services provided by Cloud Platform. It primarily helps customers securely manage access permissions to resources under the cloud platform account. With CAM, you can create, manage, and terminate users (groups), and control who can use which resources through identity and policy management. You can use CAM to bind a user or user group to a policy that allows or denies them access to specified resources to complete specified tasks. For more basic information about CAM policies, please see Policy Syntax. For more use information on CAM policies, refer to Policy. If you do not need to manage the cloud access to CVM-related resources for your sub-account, you can skip this section. Skipping these sections does not affect your understanding and use of the rest of the documentation. This feature is currently in the gray release phase, and ticket submission for application is required. Getting Started A CAM policy must authorize the use of one or more CVM operations, or it must deny the use of one or more CVM operations. You must also specify the resources that can be used for the operation (it can be all resources, or some operations can be partial resources). The policy can also contain the conditions set by the operation resources. Some API operations of CVM support resource-level permissions. That is, you cannot specify a specific resource for these API operations. Instead, you must specify all the operations resources. Authorization Policy Syntax Policy Syntax CAM policy:

{ ”version”:”2.0”, ”statement”: [{ ”effect”:”effect”, ”action”:[”action”], ”resource”:[”resource”], ”condition”:{”key”:{”value”}} }] }

The version is required, currently only allowing a value of ”2.0”. statement is used to describe the detailed information of one or more permissions. This element includes the permissions or permission sets of multiple other elements such as effect, action, resource, and condition. A policy has one and only one statement element. 1.action used to describe the operations that are allowed or denied. Operations can be APIs (described with a name prefix) or feature sets (a specific set of APIs described with a permid prefix). This element is required. 2.resource describes the detailed data authorized. Resources are described using

six paragraphs. The resource definition details will vary for each product. For information on how to specify resources, please refer to the product documentation corresponding to the resource declaration you have written. This element is required. 3.conditions describes the constraints under which the policy takes effect. Conditions consist of operators, operation keys, and operation values. Condition values can include information such as time and IP address. Some services allow you to specify additional values in the condition. This element is optional. 4.effect describes whether the statement produces an allow or explicit deny result. It includes allow and deny. This element is required. CVM Operations In the CAM policy statement, you can specify any API operation from any service that supports CAM. For CVM, please use name/cvm: as the prefix of APIs. For example: name/cvm:RunInstances or name/cvm:ResetInstancesPassword. If you want to specify multiple actions in a single statement, separate them with commas as follows:

”action”:[”name/cvm:action1”,”name/cvm:action2”]

You can also use wildcard characters to specify multiple actions. For example, you can specify all actions whose names begin with the word ” Describe ” as follows:

”action”:[”name/cvm:Describe*”]

To specify all operations in a CVM, please use a wildcard * as follows:

”action”： [”name/cvm:*”]

CVM Resource Paths Each CAM policy statement has its own resources that apply to it. The general format of resource paths is as follows:

qcs:project_id:service_type:region:account:resource

project_id: describes the project information. It is only for compatibility with early CAM logic and does not need to be filled in. service_type: the product abbreviation, such as VPC region: regional information, such as bj account: the root account information of the resource owner, such as uin/164256472 resource: specific resource details of each product, such as instance/instance_id1 or instance/* For example, you can specify an instance (i-15931881scv4) in the statement, as shown below:

”resource”:[”qcs::cvm:bj:uin/164256472:instance/i-15931881scv4”]

You can also use the * wildcard character to specify all instances belonging to a specific account, as follows:

”resource”:[”qcs::cvm:bj:uin/164256472:instance/*”]

If you want to specify all resources, or if a specific API operation does not support resource-level permissions, use the * wildcard in the Resource element as shown:

"resource":["*"]

To specify multiple resources in one instruction, please separate them with a comma. The following is an example of specifying two resources:

"resource":["resource1","resource2"]

The following table describes the resources that can be used by CVM and the corresponding methods of describing these resources. In the following table, words prefixed with $ are aliases.Among them, project refers to the project ID.Region refers to the region.Account refers to the account ID. Conditional Key for CVM In a policy statement, you can optionally specify conditions that control when a policy takes effect. Each condition consists of one or more key-value pairs. Condition keys are not case sensitive. If you specify multiple conditions or multiple keys in a single condition, we evaluate them using a logical AND operation. If you specify a key with multiple values in a single condition, we evaluate them using a logical OR operation. All conditions must be matched for permission to be granted. The following table describes the service-specific condition keys that CVM uses.

Condition Key	Reference Type	Key-value Pair
cvm:instance_type	String	cvm:instance_type=instance_type ●Among them, instance_type refers to the instance type (for example, S1.SMALL1).
cvm:image_type	String	cvm:image_type=image_type ●Image_type refers to the image type (for example, IMAGE_PUBLIC).
vpc:region	String	vpc:region=region ●The region refers to the region (for example, ap-guangzhou).
cvm:disk_size	Integer	cvm:disk_size=disk_size ●Disk_size refers to the disk size (e.g. 500).
cvm:disk_type	String	cvm_disk_type=disk_type ●Disk_type refers to the disk type (e.g. CLOUD_BASIC).
cvm:region	String	cvm:region=region ●The region refers to the region (for example, ap-guangzhou).

CVM Introduction

Overview

Last Updated At: 2025-08-04 17:12:41

Cloud Virtual Machine (CVM) provides you with secure and reliable elastic computing services. Using CVMs eliminates the need to estimate resource usage and make upfront investments required with traditional servers, helping you quickly launch any number of VMs (dependent on physical resources) in a short time frame and instantly deploy applications. CVM supports users to customize all resources: CPU, memory, hard disk, network, security, etc., and can easily adjust them when they need change.

Relevant Concepts

Last Updated At: 2025-08-04 17:12:41

Before you use CVM, you also need to understand the following concepts: ●Instance: Virtual computing resources in the cloud, including the most basic computing components such as CPU, operating system, network, and disk. ●Instance Type: Various CPU, memory, storage, and network configurations of CVMs. ●Image: Refers to the pre-made templates that run on the CVM, including pre-configured operating systems and pre-installed software. CVM provides a variety of pre-made images such as Windows and Linux. ●Local Disk: A device on the same physical server as an instance that can be used as persistent storage by the instance. ●CBS: The distributed persistent block storage device provided can be used as the system disk or expandable data disk of the instance. ●VPC: A virtual, isolated network space that is logically isolated from other resources. ●IP Address: Private IP and Public IP. Simply put, the private IP provides local area network (LAN) services, allowing CVMs to access each other within the LAN. Public IP allows users to access Internet services on CVM instances. ●Elastic IP: Static public IP designed specifically for dynamic networks to meet the needs of rapid troubleshooting. ●Security Group: A security group can be understood as a virtual firewall with status detection and packet filtering features. It is used for network access control of one or more CVMs. The security group is an important means of network security isolation. ●Log-in Methods: High-security SSH Key Pair and standard Log-in Password with regular password. ●Regions and Availability Zones: The location where instances and other resources are launched. ●Cloud Console: Web-based user interface.

Related Products

Last Updated At: 2025-08-04 17:12:41

●You can use auto-scaling to automatically increase or decrease the number of server clusters based on timing or conditions. For more information, see Elastic Scaling Product Documentation. ●You can use Cloud Load Balancer (CLB) to automatically distribute request traffic from clients across multiple CVM instances. For more information, please see Cloud Load Balancer Product Documentation. ●You can use Kubernetes Engine to manage the application lifecycle on a group of CVMs. For more information, see Product Documentation. ●You can use the Cloud Monitor service to monitor CVM instances and their system disks. For more information, see Cloud Monitoring Product Documentation.

Using CVM instances

Last Updated At: 2025-08-04 17:12:41

The following methods are provided for configuration and management of cloud servers: ●Console: Provides a web service interface for configuring and managing CVMs. ●API: Provides API to facilitate your management of CVM. For API description, see API Overview. ●SDK: You can use SDK programming (supporting PHP/Python/Java/.NET/Node.js) to invoke CVM APIs.

xx

Regions and AZs

Region

Last Updated At: 2025-08-12 20:35:22

Region refers to the geographical area of a physical data center. Cloud PlatformDifferent regions are completely isolated from each other to ensure maximum stability and fault tolerance between regions. To reduce access latency and improve download speed, it is recommended to choose the region closest to your customers. You can query the list of regions through the API to see the complete list of regions.
Related Features ●Networks in different regions are completely isolated, and cloud services in different regions cannot communicate through the private network by default. ●Cloud services in different regions can access the Internet through public network IP. Cloud services in VPC can also communicate via Peering Connection provided by Cloud Platform over the Cloud Platform NSFCNET for a more stable and faster interconnection than Internet access. ●Cloud Load Balancer (CLB) currently defaults to supporting intra-region traffic forwarding and binding with CVMs within the same local region. If the cross-region binding feature is enabled, it enables the CLB to bind CVMs across regions.

Availability Zone Related Features

Last Updated At: 2025-08-05 10:54:32

• Basic network servers in different availability zones within the same region can communicate via the internal network. • In the same region, across different availability zones, cloud products under the same VPC can interconnect through the internal network, directly accessing each other using internal IP addresses. • The aforementioned internal network interconnection applies to resources under the same account, whereas resources belonging to different accounts are completely isolated on their respective internal networks.

Availability Zone

Last Updated At: 2025-08-05 10:54:39

Availability Zone refers to Cloud Platform physical data centers within the same region with an independent power supply and network. The goal is to ensure that failures within the availability zones are isolated from each other (except for major disasters or power failures) preventing fault propagation and ensuring that users' businesses continue to provide online services. By launching instances in independent availability zones, users can protect applications from the impact of single-location failures. You can query the list of availability zones through the API to view the complete availability zone list.

How to Choose Regions and Availability Zones?

Last Updated At: 2025-08-04 17:16:12

When choosing regions and availability zones, you need to consider the following factors: ●The region where your CVM is located, and the geographical location of you and your target users. It is recommended that you choose the region closest to your customers when purchasing CVMs, to reduce access latency and improve access speed. ●Relationship between CVMs and other cloud services. It is recommended that you choose other cloud services with the same availability zone in the same region. Then the cloud products can communicate through the private network to reduce access latency and improve access speed. ●Considerations on business high availability and disaster recovery. Even if there is only one VPC, it is recommended that you deploy services at least in different availability zones to ensure fault isolation between availability zones and enable cross-availability zone disaster recovery. ●Communications between different availability zones may have network communication delays. This needs to be evaluated in combination with the actual needs of services to find the best balance between high availability and low latency. ●If you need to access hosts in other countries and districts, it is recommended that you select CVMs in other countries and districts. If you create a CVM in China, accessing hosts in other countries and districts will have high access latency. It is not recommended for such usage.

Resource Position Description

Last Updated At: 2025-08-12 20:35:22

This document explains Cloud Platform which resources are global, which are region-specific but not distinguished by availability zones, and which are based on availability zones.

Resources	Resource ID Format : 8-digit numbers and characters.	Type	Note
User account	Unlimited.	Unique globally.	Users can use the same account to access Cloud Platform resources around the world.
SSH Key	skey-xxxxxxxx	Available for all regions.	Users can use the SSH key to bind CVMs in any region under their account.
CVM Instance	ins-xxxxxxxx	Can only be used in a single availability zone of a single region.	Users can only create CVM instances in availability zones.
Custom Image	img-xxxxxxxx	Available for multiple availability zones in a single region.	Users can create custom images of instances and use them under different availability zones in the same region. If you need to use it in other regions, copy the custom image to another region by using the copy image feature.
Elastic IP	eip-xxxxxxxx	Available for multiple availability zones in a single region.	EIP addresses are created in a certain region, and can only be associated with instances from the same region.

Resources	Resource ID Format : 8–digit numbers and characters.	Type	Note
Security Group	sg–xxxxxxxx	Available for multiple availability zones in a single region.	Security groups are created in a certain region, and can only be associated with instances from the same region. Cloud Platform automatically creates three default security groups for the user.
Cloud Block Storage	disk–xxxxxxxx	Can only be used in a single availability zone of a single region.	Users can only create CBS under a specific availability zone and mount them on instances of the same availability zone.
Snapshot	snap–xxxxxxxx	Available for multiple availability zones in a single region.	After a snapshot is created for a certain CBS, users can use the snapshot to perform other operations (such as creating CBS) in this region.
CLB	clb–xxxxxxxx	Available for multiple availability zones in a single region.	The CLB can bind CVMs in different availability zones of a single region for traffic forwarding.
Virtual Private Cloud	vpc–xxxxxxxx	Available for multiple availability zones in a single region.	A VPC is created in a certain region, and resources belonging to the same VPC can be created in different availability zones.

Resources	Resource ID Format : 8–digit numbers and characters.	Type	Note
Subnet	subnet–xxxxxxxx	Can only be used in a single availability zone of a single region.	Text
Route Table	rtb–xxxxxxxx	Available for multiple availability zones in a single region.	The user needs to specify a specific VPC when creating a routing table. Thus it adheres to the position attributes of the specified VPC.

Related Operations

Migrating Instances to Other Availability Zones

Last Updated At: 2025-08-05 10:54:32

An instance that has been started cannot change its availability zone, but users can migrate the instance to another availability zone in other ways. The migration process includes creating a custom image from the original instance, starting the instance in the new availability zone using the custom image, and updating the configuration of the new instance. 1.Create a custom image of the current instance. For more information, see Create a Custom Image. 2.If the network environment of the current instance is a VPC and the current private IP address needs to be retained after migration, users can first delete the subnet in the current availability zone and then create a subnet in the new availability zone with the same IP address range as the original subnet. Note that only subnets that do not contain available instances can be deleted. Therefore, all instances in the current subnet should be moved to the new subnet. 3.Create an instance in the new availability zone using the custom image you just created. The user can select the same instance type and configuration as the original instance or select a new instance type and configuration. For more information, see Purchase and Launch an Instance. 4.If the original instance has been associated with an EIP address, it is disassociated from the old instance and associated with the new instance. For more information, see Elastic IP.

Quick Start

xx

Last Updated At: 2025-08-13 16:51:43

How to Get Started with CVM

How to Get Started with CVM

Last Updated At: 2025-08-13 16:09:36

To help you use Cloud Virtual Machine (CVM) effectively, this document provides a guide to getting started with CVM. Getting Started with CVM The Getting Started section helps you understand the basic concepts of CVM. It is suitable for users who have no basic knowledge or are new to Cloud Platform services. You can learn the following points: What is CVM? The Features and Strength of CVM Advanced Edition The Advanced Edition helps you choose the right cloud server for you when applying: Before using the cloud server, you need to complete registration and verification first. If you're unsure how to select configurations, we provide configuration recommendations. Before purchasing and using a CVM, we will help you make a better choice that suits your needs: When you are faced with a variety of models and do not know how to choose, CVM Model Selection helps you understand the applicable scenarios and performance of different models, and helps you choose the model that suits your business scenario. If you are not sure where to configure, Region and Availability Zone helps you understand the best options for regions and availability zones. Hands-On Guide The Hands-On Guide provides detailed operational instructions to help you quickly get started with simple usage of Windows and Linux system cloud servers, guiding you through account registration, login, management, and other steps. General Approach 1.Register Account 2.Determine Region and Cloud Server Configuration 3.Create Cloud Server 4.Log in to Cloud Server 5.Format Data Disk and Perform Partition Operations 6.Install and Set Up Application Environment For detailed operational instructions, refer to Quick Start Guide for Windows Cloud Server and Quick Start Guide for Linux Cloud Server. Advanced Guide The Advanced Guide offers more detailed operational instructions for managing cloud servers, assisting you in environment setup and software installation among others. By following this guide, you will accomplish operations and maintenance deployment for Windows and Linux system cloud servers. Windows System Cloud Server Operations Manual 1.Log in to Windows Cloud Server 2.Partition and Format Data Disk 3.Environment Setup 4.System Maintenance Linux System Cloud Server Operations Manual 1.Log in to Linux Cloud Server 2.Mount Data Disk on Linux 3.Install Software 4.Environment Setup 5.Upload Files 6.Common Linux Operations and Commands Others Adjust instance configuration: If you choose a lower hardware configuration when your application is in its early stages and the number of requests is small, and as your application grows rapidly and the number of service requests increases dramatically, you can adjust the instance configuration by quickly upgrading the hardware to improve service processing speed and better meet your changing needs. FAQs: If you still encounter other FAQs of CVM Management, we provide a collection of frequently asked questions for your reference, so that you can quickly locate and solve your questions. Issue Feedback: If you have any unresolved questions, feel free to let us know! We'll be quick to address your concerns and clarify things for you.

Quick Start Guide for LINUX Server

Step One: Preparation and Selection

Last Updated At: 2025-08-13 16:09:36

Register for a Cloud Platform Account New users need to register on the official website of the cloud platform. Registration guidelines can be found in How to Register for the Cloud Platform. Determine the Location and Availability Zone of the Cloud Server Principles for choosing a region: -Proximity principle: Choose the region based on your users' geographical location. The closer the cloud server is to your visitors, the lower the access latency and the higher the download speed. -Intranet communication same region principle: Within the same region, intranet is interconnected; in different regions, intranet does not connect. Users who need multiple cloud servers to communicate via intranet must choose the same server region. Cloud servers in the same region can communicate with each other via intranet (intranet communication). Cloud servers in different regions cannot communicate with each other via intranet (communication requires the public network). Determine the Configuration Plan for the Cloud Server The cloud platform provides the following recommended configurations: -Entry-level [Recommended Selection]: Suitable for personal websites at the start-up stage. For example: small websites such as personal blogs. -Basic Type: Suitable for websites or applications with certain traffic volume. For example: larger corporate official websites, small e-commerce websites. -Popular Type: Suitable for demands that frequently use cloud computing and have some computational requirements. For example: portal websites, SaaS software, small apps. -Application Type: Suitable for applications with higher concurrency requirements and scenarios where there are certain requirements for the network and computing performance of the cloud server. For example: large portals, e-commerce websites, game apps. If the recommended configuration does not meet your needs, you can compare various configuration plans according to actual needs in "More Models". Of course, after applying for a cloud server, you can upgrade or downgrade the configuration at any time based on your needs.

Quick Start Guide for LINUX Server

Last Updated At: 2025-08-13 16:09:36

This document mainly introduces how to quickly use the functions related to Linux system cloud server instances, guiding beginners to quickly understand the creation and configuration of cloud servers.

Step Two: Create a Linux Cloud Server

Last Updated At: 2025-08-13 16:17:51

This step introduces how to create a Linux cloud server. The cloud platform offers both quick setup and custom setup options. This section uses quick setup as an example. If quick setup doesn't suit your needs, refer to the document for configuring a Linux cloud server yourself. 1.Log in to the official website of the cloud platform, select [Cloud Products] – [Computing & Networking] – [Cloud Server], click the [Buy Now] button to enter the cloud server application page. 2.Choose the image. 3.Select the model. 4.Choose the region. Choosing a region closer to your customers reduces access latency and increases download speed. 5.Select public network bandwidth. If no connection to the public network is needed, set the bandwidth value to 0. 6.Select the number of servers and the duration of the application. Quick setup uses automatically generated passwords. After creation, the password will be sent to you through internal mail. To view more default configurations, hover over [More Default Configurations] at the top of the quick application page. See the next steps for viewing internal mail.

Step Three: Login to the Linux Cloud Server

Last Updated At: 2025-08-13 16:17:51

This section describes common methods for logging into a Linux cloud server. Different login methods can be used under different circumstances. Here we introduce console login. For more login methods, see Logging into Linux Instances. Prerequisites When logging into a cloud server, you need to use the administrator account and corresponding password. • Administrator Account: For instances of the Linux type, the administrator account is uniformly root. • Password: In quick setup, the initial password is randomly assigned by the system. Specific viewing operations are found in the next step (viewing internal mail and cloud server information). Console Login to the Cloud Server Enter the account (root) and initial password (or the password you modified) to log in. Note: This terminal is exclusive, meaning only one user can use console login at any given time.

Step Four: Partitioning and Formatting the Data Disk

Last Updated At: 2025-08-13 16:17:51

Prerequisites • For users who have applied for a data disk, it needs to be formatted before it can be used. Users without a data disk can skip this step. • Ensure that you have completed the operations in Step Three and logged into the cloud server. • For hard disks larger than 2TB, please use GPT method to mount the data disk operation. Details can be found in Using GPT Partition Table to Partition and Format. Partitioning the Data Disk 1.Log into the Linux cloud server using the method described in Step Three. Note: Only data disks support partitioning; system disks do not. If you forcefully partition the system disk, it may lead to serious issues such as system crashes. In such cases, the cloud platform does not assume liability for compensation. 2.Enter the command fdisk –l to view your data disk information. In this example, there is a 54GB data disk (/vdb) that needs to be mounted. Note: Both fdisk –l and df –h are commands used to check data disk information, but before partitioning and formatting the data disk, using the df–h command will not show the data disk. 3.Partition the data disk. Follow the prompts on the interface sequentially: 1.Enter fdisk /dev/vdb (to partition the data disk), press enter; 2.Enter n (create a new partition), press enter; 3.Enter p (create an extended partition), press enter; 4.Enter 1 (use the first primary partition), press enter; 5.Press enter (use default settings); 6.Press enter again (use default settings); 7.Enter wq (save the partition table), press enter to start partitioning. In this example, we create one partition, but developers can also create multiple partitions based on their needs. 4.Use the fdisk–l command to see that the new partition vdb1 has been successfully created. Formatting the Data Disk 1.After formatting a new partition, you need to format the well–divided areas. You can decide on the format of the file system yourself, such as ext2, ext3, etc. This example uses ext3 as an instance. Use the following command to format the new partition: mkfs.ext3 /dev/vdb1 2.Mounting the partition: Use the following commands to create the mydata directory and mount the partition under this directory: 3.mkdir /mydata mount /dev/vdb1 /mydata Use the command to check the mounting: df —h If the selected vdb1 information appears as shown in the figure, it means the mounting was successful, i.e., you can see the data disk. 4.Setting up auto–mount at boot: If you want the cloud server to automatically mount the data disk when restarting or powering on, you must add the partition information to /etc/fstab. Use the following command to add partition information:

Use the following command to view:

By now, you have completed the creation and basic configuration of the Linux system–based cloud server.

Quick Start Guide for Windows Server

Step One: Preparation and Selection

Last Updated At: 2025-08-13 16:27:05

Register a Cloud Platform Account New users need to register on the official website of the cloud platform. Registration guidance can be found in How to Register for the Cloud Platform. Determine the Region and Availability Zone for Your Cloud Server Principles for Choosing a Region: -Closest to Users Principle: Please select the region for your cloud server based on the geographical location of your users. The closer the cloud server is to its visitors, the smaller the latency and higher the access speed you will achieve. -Intranet Communication Same Region Principle: Within the same region, intranet communication is possible; between different regions, intranet communication is not available. Users requiring multiple cloud servers to communicate via the intranet must choose the same region for their cloud servers. Servers within the same region can communicate through the intranet (intranet communication). Servers located in different regions cannot communicate through the intranet (communication requires passing through the public network). Determine the Configuration Plan for Your Cloud Server The cloud platform offers the following recommended configurations: [Recommended Models]-Entry Level: Suitable for personal websites in their early stages. For example: Personal blogs and other small-scale websites. -Basic Model: Appropriate for websites or applications with a certain level of traffic. For example: Larger corporate homepages, small e-commerce sites. -General Model: Suitable for frequent use of cloud computing services and demands of moderate computational load. For example: Portal websites, SaaS software, small mobile apps. -Application-Oriented Model: Applicable to high-concurrency requirements and scenarios where there are specific needs for the networking and computational performance of the cloud server. For example: Large portals, e-commerce platforms, gaming apps. If the recommended configurations do not meet your needs, you can compare various plans according to your actual requirements under [More Models]. Of course, you can also upgrade or downgrade your configuration at any time after applying for the cloud server based on your needs. Note: Windows cloud servers cannot be used as public network gateways. Users requiring a public network gateway should refer to the Quick Start Guide for Linux Cloud Servers.

Step Two: Creating a Windows Cloud Server

Last Updated At: 2025-08-13 16:27:05

This step describes how to create a Windows cloud server. The cloud platform offers both quick setup and custom configuration options. This section uses quick setup as an example; if quick setup does not meet your requirements, you can consult the Custom Configure Windows Cloud Server documentation for configuration. 1.Log in to the official website of the cloud platform, select [Cloud Products]–[Computing & Networking]–[Cloud Server], click the [Purchase Now] button to enter the application page for cloud servers. 2.Choose an image. Select a Windows operating system that meets your needs. 3.Select a model. 4.Choose a region. Regions closer to your customers reduce access latency and improve download speeds. 5.Select public network bandwidth. If you do not need to connect to the public network, set the bandwidth value to 0. 6.Choose the number of servers and application duration. Quick setup uses an automatically generated password. After creation, the password will be sent to you via internal messaging. For more default configurations, hover your mouse over [More Default Configurations] at the top of the quick application page. To view internal messages, see the next step.

Step Three: Logging into Windows Cloud Server

Last Updated At: 2025-08-13 16:27:05

This section introduces common methods for logging into a Windows cloud server. Different login methods can be used under different circumstances. Here we will introduce console login; for more login methods, please refer to Log into Windows Instance. Prerequisites When logging into the cloud server, you need to use the administrator account and its corresponding password. • Administrator Account: For instances of the Windows type, the administrator account is uniformly set as Administrator. • Password: In quick configuration, the initial password is randomly assigned by the system. In the next step (viewing internal messages and cloud server information), specific viewing operations will be detailed. For more details, please see Login Password. View Internal Messages and Cloud Server Information After completing the application and creation of the cloud server, instance name, public IP address, private IP address, login name, initial login password, and other information will all be sent to your account via internal messaging. 1.Log into the cloud server management console. After logging in, you will be able to see the public IP address, private IP address, and other information. 2.Click on [Internal Messaging] in the upper right corner. 3.On the internal messaging page, you can view newly created cloud servers along with login names, passwords, and other information.

Step Four: Formatting and Partitioning Data Disks

Last Updated At: 2025-08-13 16:27:05

Here we will illustrate formatting using Windows 2012 R2 as an example. Prerequisites • Users who have applied for data disks need to format them before they can be used. Users who have not applied for data disks may skip this step. • Ensure that you have completed the operations described in Step Three, having logged into the cloud server. Formatting the Data Disk 1.Log into the Windows cloud server using the method introduced in Step Three. 2.Click [Start] – [Server Manager] – [Tools] – [Computer Management] – [Storage] – [Disk Management].

4.Right-click again and choose [Initialize Disk]: 5.Depending on the partition style, select either [GPT] or [MBR], then click [OK]: Note: When the disk size exceeds 2TB, only GPT partitioning is supported. If you are unsure whether the disk will be expanded beyond this value in the future, it is recommended that you choose GPT partitioning. If you are certain that the disk size will not exceed this value, MBR partitioning is suggested for better compatibility. Disk Partitioning (Optional) 1.Right-click on the unallocated space and select [New Simple Volume]: 2.In the pop-up window titled ”New Simple Volume Wizard,” click [Next]: 3.Enter the required disk size for the partition, then click [Next]: 4.Assign a drive letter, then click [Next]: 5.Choose the file system and format the partition, then click [Next]: 6.Complete creating the simple volume by clicking [Finish]: 7.Open [My Computer] from the [Start] menu to check the new partition: Up to this point, you have successfully completed the creation and basic configuration of the Windows-based cloud server.

Quick Start Guide for Windows Server

Last Updated At: 2025-08-14 09:27:06

This document primarily introduces how to quickly utilize the features related to instances of Windows system cloud servers, guiding newcomers to swiftly understand the creation and configuration of cloud servers.

Network Planning

Network Planning

Last Updated At: 2025-08-13 16:27:05

The Virtual Private Cloud (VPC) is a logically isolated network space customized by users on the cloud platform. Within the VPC, users can freely define subnet segmentation, IP addresses, routing policies, etc., so the cloud platform recommends choosing a VPC for users. In order for users to better utilize the VPC, the cloud platform provides the following suggestions regarding network planning:

Determine the Number of VPCs

Last Updated At: 2025-08-13 16:51:43

Known Features: o A VPC has regional attributes. By default, the intranet between cloud service products in different regions is not interconnected. When cross-regional communication is needed, inter-region connectivity can be achieved through establishing peer connections. o By default, the intranets within different VPCs in the same region are not interconnected. When cross-VPC communication is needed, inter-VPC connectivity can be established through setting up peer connections. o By default, the intranets between different availability zones within the same VPC are interconnected. • Related Suggestions: o When your business requires multi-regional deployment systems, multiple VPCs are inevitably necessary; you can choose to establish VPCs in regions closer to customers to reduce latency and improve access speed. o When you have multiple sets of businesses deployed in the current region and hope to isolate networks among different businesses, you can create respective VPCs for each business in the current region. o When you do not have multi-regional deployment needs nor any requirements for network isolation among various businesses, you can simply use one VPC.

Determine Subnet Segmentation

Last Updated At: 2025-08-13 16:51:43

Known Features: o A subnet is an IP address block within the VPC, and all cloud resources in the VPC must be deployed within subnets. o Under the same VPC, subnet segments cannot overlap. o Currently, VPC supports three segments of intranet IPs: 10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16. o Once a VPC is successfully created, the segment cannot be modified. • Related Suggestions: o If only subnet planning for the VPC is involved without consideration for communication with the base network or IDC network, you can choose any of the above segments to create a new subnet. o If a VPN connection is needed, the local segment (VPC segment) and remote segment (your IDC segment) must not overlap, so when creating a new subnet, avoid overlapping with the remote segment. o When dividing network segments, consider the capacity of IP addresses in the segment, i.e., how many available IP addresses there are. o Finally, it's recommended to divide subnets according to business modules within the same VPC, such as subnet A for the web layer, subnet B for the logic layer, and subnet C for the DB layer, which facilitates access control and filtering combined with Network Access Control Lists (NACL).

Determine Routing Policies

Last Updated At: 2025-08-13 16:51:43

Known Features: o The route table consists of a series of routing rules that control the direction of outbound traffic from subnets within the private network (VPC). o Each subnet must be associated with a route table and can only be associated with one route table. o One route table can be associated with multiple subnets. o When a user creates a private network, the system automatically generates a default route table for them, meaning that the intranet within the private network is interconnected. • Related Suggestions: o If there's no need for special control over the flow direction of traffic from subnets, under the condition of default VPC intranet interconnection, the default route table can be used without configuring custom routing strategies; o If special control over the flow direction of traffic from subnets is needed, you can refer to the official website for detailed instructions on using route tables. For more about VPC introductions, please refer to Private Network.

Selecting Cloud Hard Drives

Selecting Cloud Hard Drives

Last Updated At: 2025-08-13 16:51:43

To meet the diverse needs of different customers across various scenarios, the cloud platform offers the following recommendations for selecting hard drives based on application scenarios: Scenarios for SSD Local Disks • Low Latency: Provides microsecond-level access delay. • Large Online Application Logs: Large online applications generate vast amounts of log data requiring high-performance storage, where the reliability requirement for stored logs is relatively low. • Used as Temporary Read Cache: Local SSDs excel in random read performance (4KB/8KB/16KB random reads), making them suitable for read-only secondary databases of relational databases like MySQL, Oracle, etc. Due to the higher cost of memory compared to solid-state disks, local SSDs can also serve as second-tier caches for caching-oriented services like Redis, Memcache. • Single Point of Failure Risk: There is a risk of single-point failure; it's recommended to ensure data availability through redundancy at the application layer. For core business scenarios, SSD cloud disks are suggested. Scenarios for Standard Cloud Disks • Cost-effective storage with the same data durability as SSD cloud disks, suitable for cold data backup and archiving of critical businesses, with a single disk capacity reaching up to 16TB. • Applicable for large file sequential read/write scenarios such as log streams, streaming media services, data warehousing, etc., meeting offline analysis needs for TB-level massive datasets under the Hadoop framework. • Not suitable for supporting OLTP (Online Transaction Processing) core business operations. Scenarios for SSD Cloud Disks • High Performance, High Data Reliability: Utilizes industry-leading NVMe solid-state storage as the disk medium. Suitable for I/O-intensive applications while providing long-term stable ultra-high performance per single disk. • Medium-to-Large Database Applications: Can support mid-to-large-scale relational database applications like MySQL, Oracle, SQL Server, MongoDB, capable of handling tables with millions of rows. • Core Business Systems: Suitable for I/O-intensive core business systems requiring high data reliability. • Big Data Analysis: Provides distributed processing capabilities for TB, PB level datasets, applicable in areas such as data analysis, mining, business intelligence. Select Instance Type In order to meet the diverse needs of different customers for various application scenarios, the cloud platform offers the following instance type selection recommendations under different application scenarios: • Personal Website It is recommended to use Standard instances, suitable for general workloads such as small-to-medium-sized web applications and databases. • Corporate Websites / E-commerce / Apps Standard instances are also recommended, suitable for general workloads including small-to-medium-sized web applications and databases. • Relational Database / Distributed Cache Memory-intensive instances are recommended, suitable for applications requiring extensive memory operations, searches, and calculations. • NoSQL Databases High IO instances are recommended, suitable for I/O intensive scenarios where high disk read/write speed and low latency are required, such as NoSQL

databases (e.g., MongoDB), clustered databases, etc. • High-performance Computing Compute-intensive instances or compute network-enhanced instances are recommended, suitable for applications consuming significant computing resources, such as large-scale online games, high-performance engineering scientific applications, video encoding, etc. • High-performance Online Games Compute-intensive instances or compute network-enhanced instances are recommended, suitable for applications consuming significant computing resources, such as large-scale online games, high-performance engineering scientific applications, video encoding, etc. • Mobile Games / Web Games Compute-intensive instances or compute network-enhanced instances are recommended, suitable for applications consuming significant computing resources, such as large-scale online games, high-performance engineering scientific applications, video encoding, etc. • Live Streaming Standard network-enhanced or compute network-enhanced instances are recommended, paired with 25G NICs, offering improved network performance by 2.5 times compared to ordinary 10G data centers; featuring higher bandwidth and lower latency. • Finance Dedicated Host Standard instances are recommended, unlike regular standard instances, they offer exclusive physical servers with resource isolation; enhanced security control allowing users to define their own server specifications; meeting strict regulatory requirements in the financial industry. • Scientific Computing GPU Compute instances are recommended, suitable for deep learning, scientific computing such as computational fluid dynamics, computational finance, genomics research, environmental analysis, high-performance computing, and other GPU computing workloads on the server side. • Machine Learning GPU Compute instances are recommended, suitable for deep learning, scientific computing such as computational fluid dynamics, computational finance, genomics research, environmental analysis, high-performance computing, and other GPU computing workloads on the server side. • Rendering GPU Rendering instances are recommended, used for non-linear editing, video transcoding, graphics acceleration visualization, and 3D design in GPU rendering scenarios. • Hadoop / Spark / ElasticSearch Big Data instances are recommended, suitable for distributed computing environments such as Hadoop (HDFS/MapReduce/Spark/Hive, etc.), massively parallel processing (MPP) data warehouses, log or data processing applications, etc. For more application scenarios, please refer to the Introduction to Instance Types. Configure Security Groups Security groups are instance-level firewalls provided by cloud platforms, enabling inbound/outbound traffic control for any cloud server. 1.Log in to the cloud server console and click [Security Groups] in the left navigation pane. 2.Click the [New] button, enter a name for the security group (for example, my-security-group), choose to create using a template or custom creation, confirm ingress and egress rules, then click [OK]. 3.On the right side of the security group list, click the [Add Instances] button, check the required associated virtual machines, thus completing the operation of associating VMs with the security group. Alternatively, You may also go to the VM list page to view or modify the security groups bound to a certain VM. In the [VM] list page, select the VM whose security group needs adjustment, then click [More]>[Configure Security Group] on the right, and choose the security group binding.

Select Instance Type

Last Updated At: 2025-08-11 14:58:27

In order to meet the diverse needs of different customers for various application scenarios, the cloud platform offers the following instance type selection recommendations under different application scenarios: • Personal Website It is recommended to use Standard instances, suitable for general workloads such as small-to-medium-sized web applications and databases. • Corporate Websites / E-commerce / Apps Standard instances are also recommended, suitable for general workloads including small-to-medium-sized web applications and databases. • Relational Database / Distributed Cache Memory-intensive instances are recommended, suitable for applications requiring extensive memory operations, searches, and calculations. • NoSQL Databases High IO instances are recommended, suitable for I/O intensive scenarios where high disk read/write speed and low latency are required, such as NoSQL databases (e.g., MongoDB), clustered databases, etc. • High-performance Computing Compute-intensive instances or compute network-enhanced instances are recommended, suitable for applications consuming significant computing resources, such as large-scale online games, high-performance engineering scientific applications, video encoding, etc. • High-performance Online Games Compute-intensive instances or compute network-enhanced instances are recommended, suitable for applications consuming significant computing resources, such as large-scale online games, high-performance engineering scientific applications, video encoding, etc. • Mobile Games / Web Games Compute-intensive instances or compute network-enhanced instances are recommended, suitable for applications consuming significant computing resources, such as large-scale online games, high-performance engineering scientific applications, video encoding, etc. • Live Streaming Standard network-enhanced or compute network-enhanced instances are recommended, paired with 25G NICs, offering improved network performance by 2.5 times compared to ordinary 10G data centers; featuring higher bandwidth and lower latency. • Finance Dedicated Host Standard instances are recommended, unlike regular standard instances, they offer exclusive physical servers with resource isolation; enhanced security control allowing users to define their own server specifications; meeting strict regulatory requirements in the financial industry. • Scientific Computing GPU Compute instances are recommended, suitable for deep learning, scientific computing such as computational fluid dynamics, computational finance, genomics research, environmental analysis, high-performance computing, and other GPU computing workloads on the server side. • Machine Learning GPU Compute instances are recommended, suitable for deep learning, scientific computing such as computational fluid dynamics, computational finance, genomics research, environmental analysis, high-performance computing, and other GPU computing workloads on the server side. • Rendering GPU Rendering instances are recommended, used for non-linear editing, video transcoding, graphics acceleration visualization, and 3D design in GPU rendering scenarios. • Hadoop / Spark / ElasticSearch Big Data instances are recommended, suitable for distributed computing environments such as Hadoop (HDFS/MapReduce/Spark/Hive, etc.), massively parallel processing (MPP)

data warehouses, log or data processing applications, etc. For more application scenarios, please refer to the Introduction to Instance Types.

Configure Security Groups

Last Updated At: 2025-08-11 14:58:27

Security groups are instance-level firewalls provided by cloud platforms, enabling inbound/outbound traffic control for any cloud server. 1.Log in to the cloud server console and click [Security Groups] in the left navigation pane. 2.Click the [New] button, enter a name for the security group (for example, my-security-group), choose to create using a template or custom creation, confirm ingress and egress rules, then click [OK]. 3.On the right side of the security group list, click the [Add Instances] button, check the required associated virtual machines, thus completing the operation of associating VMs with the security group. Alternatively, You may also go to the VM list page to view or modify the security groups bound to a certain VM. In the [VM] list page, select the VM whose security group needs adjustment, then click [More]>[Configure Security Group] on the right, and choose the security group binding.

Operation Guide
Cloud Block Storage
Cloud Block Storage

Last Updated At: 2025-08-14 09:25:05

Expand CBS Overview A Cloud Block Storage (CBS) is a scalable storage device in the cloud environment. Users can expand its size at any time after creating the CBS to increase storage space without losing the original data on the CBS. After expanding the CBS, you need to extend partitions and the file system. You can allocate the additional capacity to existing partitions, or format the expanded part as independent new partitions. Note The maximum disk capacity supported by the MBR partition format is 2 TB. If your disk partition is in MBR format and you need to expand it beyond 2 TB, we suggest recreating and attaching a new data disk, utilizing the GPT partition scheme, and transferring the data to the new disk. Expand Data Disk Note If multiple CBSs of the same capacity and type are connected to your cloud server, you can refer to the operations for identifying data disks to distinguish them. After selecting the data disk you wish to expand, proceed with the expansion process as outlined below. • Expand via Cloud Server Console (Recommended) 1.Log in to the Cloud Server Console. 2.From the row corresponding to the target cloud server, select More > Resource Adjustment > Expand CBS. 3.In the pop-up "Expand CBS" window, select the data disk you wish to expand, and click Next. 4.In the "Adjust Capacity" step, set the target capacity (which must be greater than or equal to the current capacity), and click Next. 5.In the "Expand Partitions and File System" step, review the precautions, and click Start Adjustment. Refer to the following illustration:

6.Depending on the operating system type of your target cloud service, you need to Extend Partitions and File System (for Windows) or Extend Partitions and File System (for Linux) to allocate the expanded portion's capacity to existing partitions, or format the expanded part as independent new partitions.

Expand System Disk 1.Log in to the Cloud Server Console, and from the row corresponding to the target cloud server, select More > Resource Adjustment > Expand CBS. 2.In the pop-up "Expand CBS" window, select the system disk you wish to expand, and click Next. 3.In the "Adjust Capacity" step, set the target capacity (which must be greater than or equal to the current capacity), and click Next. 4.To complete the expansion operation: Offline Expansion: In the "Expand Partitions and File System" step, review the precautions, check the box agreeing to force shutdown, and click Start Adjustment. Refer to the following illustration:

After completing the expansion operations on the console, check the cloudinit configurations for Linux instances or view the cloudinit configurations for Windows instances based on the actual operating systems of the corresponding cloud servers. Perform expansion operations on partitions and file systems as needed based on the confirmation results. Online Expansion: Cloud servers support online expansion of cloud disks used as system disks without interrupting services. If you wish to utilize this feature, submit an application; after approval, you may begin using it. • In the step of "Expanding Partitions and File

Systems," review the precautions and click Start Adjustment. Refer to the image below:

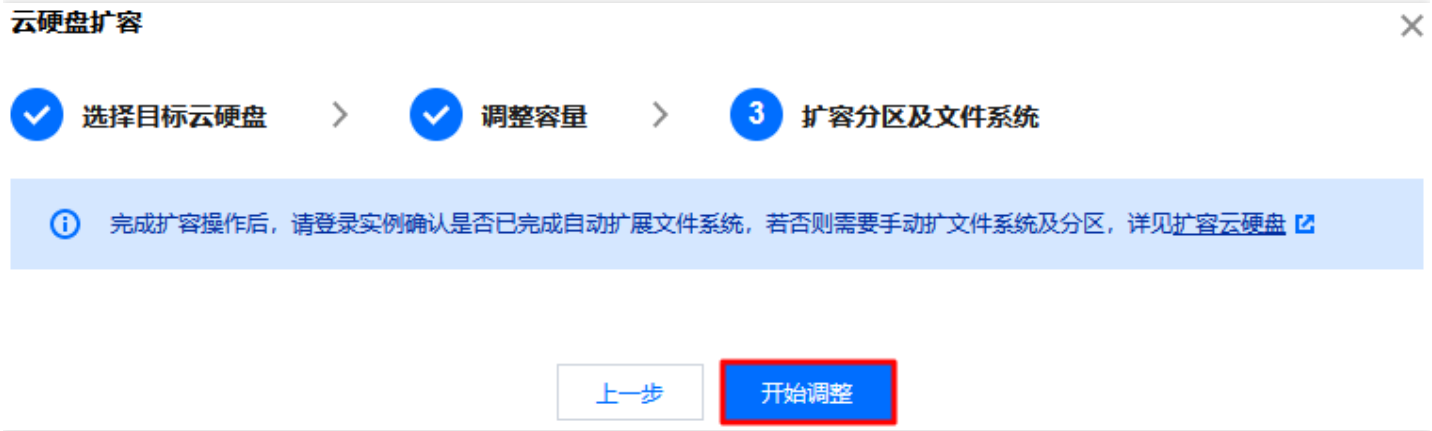

After finishing the console expansion operation, log into the instance to confirm whether the file system has been automatically extended. If not, refer to Online Expansion of System Disks and File Systems for operations on expanding partitions and file systems. Related Operations Distinguish Data Disks 1.Log in to the Linux instance. 2.Execute the following command to see the correspondence between cloud disks and device names. ls –l /dev/disk/by–id The returned results are illustrated in the image below:

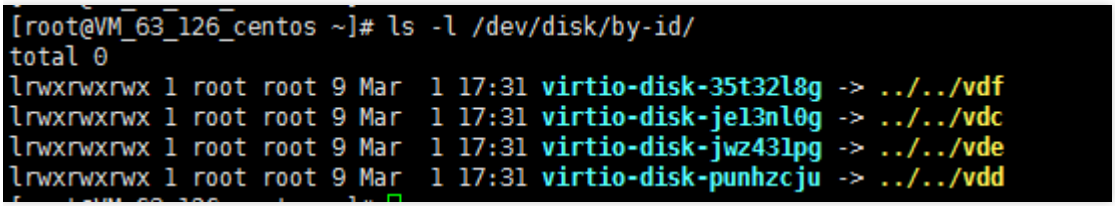

In these results, disk–xxxx represents the ID of the cloud disk; you can find it on the cloud disk management console. View Instance Cloudinit Configuration After completing expansion operations, log into the Linux instance to verify if /etc/cloud/cloud.cfg includes growpart and resizefs configuration entries. • If present, no additional steps are necessary. See the image below:

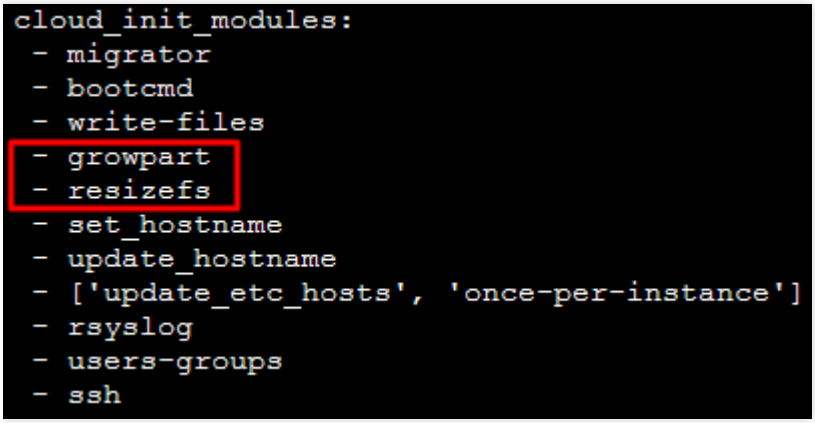

o growpart: Expands the partition size to match the disk size. o resizefs: Extends the file system of the / partition to match the partition size. • If absent, manually expand the file system and partitions based on the operating system type of the targeted cloud service. You will need to execute the procedures for expanding partitions and file systems on Linux, which involves allocating the expanded capacity to existing partitions or formatting the expanded capacity into new separate partitions. Adjusting Cloud Disk Performance Cloud disk performance generally correlates with its capacity. You can obtain higher performance by adjusting the capacity of the cloud disk when it hasn't reached its maximum performance threshold. Among them, enhanced SSD cloud disks allow breaking past baseline performance limits by configuring additional performance once they've hit their peak baseline levels. When applicable, you can configure and adjust extra performance anytime. For more details, refer to the Enhanced SSD Cloud Disk Performance Specifications. Note: • At present, only enhanced SSD cloud disks support independent performance adjustments. • Additional performance can only be adjusted independently once baseline performance reaches its maximum. • During periods of cloud disk performance adjustment, business operations and regular usage remain uninterrupted. Performance Upgrade Provided the preconditions are met, you can proceed with a performance upgrade through the following method: 1.Log in to the cloud disk control panel. 2.Choose the geographical location, then pick out the cloud disk whose performance needs enhancement. 3.Under the selected cloud disk, opt for More > Adjust Performance. 4.In the displayed "Adjust Performance" dialog box, select the desired target configuration. 5.Tick off the guidelines and commence the adjustment process. Performance Downgrade Given the prerequisites are satisfied, you can carry out a performance downgrade via the following procedure: 1.Log in to the cloud disk control panel. 2.Choose the geographical location, then pick out the cloud disk whose performance requires reduction. 3.Under the selected cloud disk, opt for More > Adjust Performance. 4.In the displayed "Adjust Performance" dialog box, select the desired target configuration. 5.Tick off the guidelines and commence the adjustment process.

Network
Elastic Network Interface Card (ENI)

Last Updated At: 2025-08-14 09:25:05

If your cloud server is required to use an Elastic Network Interface (ENI), please follow the configuration steps outlined below to complete the relevant tasks: 1.Create an elastic network interface card. Upon creation, you can view detailed information about the ENI by checking its properties. 2.Bind and configure the elastic network interface card to the cloud server. (Critical Step) 3.Configure routing tables for the cloud server and private network. 4.Allocate an internal IP address. (1) Log in to the private network console. (2) Click on the left sidebar's ENI section to enter the list page of elastic network interfaces. (3) Click on the ID/name of the elastic network interface to view detailed information about the ENI on its information page. (4) Click on IP Management to enter the detail page. (5) Click Allocate Internal IP Address, choose allocation mode (automatic assignment or manual entry—manual entry requires specifying suitable internal IP addresses), and click Confirm to finalize the operation. 5.Manage the elastic network interface. (1) Release internal IP addresses (2) Unbind from the cloud server (3) Delete the elastic network interface (4) Bind an elastic public IP address (5) Unbind the elastic public IP address (6) Modify primary internal IP address (7) Change the subnet associated with the elastic network interface

Security

Security

Last Updated At: 2025-08-14 09:25:05

Security Group Overview A security group is a stateful virtual firewall with a packet filtering function. It is used to set network access control for one or more CVMs and is an important network security isolation method provided by Cloud Platform. You can control access permission to your instances using the earlier methods: □Type:Create multiple security groups and assign different rules to each security group. □Each instance is assigned one or more security groups. These rules will be used to determine which traffic can access the instance and which resources the instance can access. □Configure the security group so that only specific IP addresses or specific security groups can access the instance. □You can modify the rules of the security group at any time, and the new rules take effect immediately. Usage Limitations: For usage limitations and quotas related to security groups, see the section on security group restrictions in the Overview of Usage Limitations. Security Group Rules: Components: Security group rules consist of the following components: • Source: The IP address of source data (inbound) or destination data (outbound). • Protocol type and port: Such as TCP, UDP, etc. • Policy: Allow or deny. Rule Priority: • Rules within a security group have priority levels. Rule priority is represented by its position in the list; the rule at the top of the list has the highest priority and is applied first; the rule at the bottom of the list has the lowest priority. • In case of conflicting rules, the default is to apply the rule located higher up in the list. • When traffic enters/exits instances bound to a certain security group, it will start matching from the topmost rule in the security group rule list until the last rule. If a match is successful for any rule (allowed through/denied), subsequent rules after this matched rule will not be checked. Multiple Security Groups An instance can bind to one or multiple security groups. When an instance binds to multiple security groups, these security groups will be executed sequentially from top to bottom, and you may adjust the priority of security groups at any time. Security Group Templates When creating a new security group, you can choose from two security group templates provided by the cloud platform: • All Ports Open Template: Will allow all inbound and outbound traffic. • Common Ports Open Template: Will allow TCP port 22 (for Linux SSH login), ports 80 and 443 (for web services), port 3389 (for Windows remote login), ICMP protocol (ping), and internal network traffic. Note: • If the provided security group templates do not meet your actual needs, you can also create custom security groups. For details, see Creating Security Groups and Security Group Application Scenarios. • If you have security protection requirements for application layer (HTTP/HTTPS), you may separately apply for the Cloud Platform's Web Application Firewall (WAF). WAF will provide you with application-layer web security protection against attacks such as web vulnerability exploitation, malicious crawlers, and CC attacks, ensuring website and web application security. Security Group Best Practices Recommendation Create Security Groups • When calling the API to apply for a CVM, it is recommended to specify a security group. If no security group is specified, the system-generated

default security group will be used. The default security group cannot be deleted, and its default rule allows all IPv4 rules. After creation, modifications can be made as needed. • In case of changes to instance protection policies, it is recommended to modify the rules within the security group first, without needing to recreate a new security group. Manage Rules • When modifying rules, current security groups can be exported and backed up first. If new rules have adverse effects, previous security group rules can be imported for recovery. • When a large number of rule entries are required, parameter templates can be utilized. Associate Security Groups • You can add instances with similar protection requirements to a single security group, without needing to configure separate security groups for each instance. • It is not recommended to bind too many security groups to a single instance, as conflicts between different security group rules may lead to network disconnections. Security Groups and Cloud Firewalls Cloud Firewall (CFW), a SaaS firewall based on public cloud environments, primarily provides internet boundary protection for users, addressing unified management and log auditing needs for cloud-based access control. Apart from traditional firewall features, Cloud Firewalls also support multi-tenancy and elastic scaling in the cloud environment, making them the first network security infrastructure for businesses migrating to the cloud. In practical scenarios, security groups are typically deployed at the boundaries of cloud products such as CVMs, enabling access control between security groups associated with these cloud products. On the other hand, Cloud Firewalls are deployed at the boundaries between VPCs or the internet, facilitating access control between VPCs or between the cloud platform and the internet. Creating Security Groups Overview A security group acts as a virtual firewall for cloud server instances; each cloud server instance must belong to at least one security group. When you create a cloud server instance and haven't created any security groups yet, the cloud platform offers two templates — "Open All Ports" and "Open ports 22, 80, 443, 3389 and ICMP protocol" — to automatically generate a default security group for you. For more details, please refer to the Security Group Overview. If you do not wish for your cloud server instance to join the default security group, you can follow the description in this article to create a custom security group yourself. This guide will walk you through creating a security group on the cloud server console. Directions 1.Log in to the Cloud Server Console. 2.In the left navigation bar, click Security Groups to enter the security group management page. 3.On the security group management page, select the region and click New. 4.In the pop-up "Create Security Group" window, complete the following configurations: • Template: Choose an appropriate template based on the services that need to be deployed by the cloud server instances within the security group to simplify the configuration of security group rules. As shown in the table below:

Template	Description	Scenario
Open All Ports	By default, all ports are open to both public and private networks, which poses certain security risks.	-

Template	Description	Scenario
Open ports 22, 80, 443, 3389 and ICMP protocol	By default, opens ports 22, 80, 443, 3389 and ICMP protocol; all internal network traffic is allowed.	The instances in the security group require deployment of web services.
Customized	After successful creation of the security group, add security group rules as needed. Specific operations can be found under Adding Security Group Rules.	–

Name: Customize the name of the security group. • Project Affiliation: Default selection is “Default Project,” but it can be specified for other projects for easier future management. • Remarks: Customizable — briefly describe the security group for easier future management. • Advanced Options: You can configure tags for the security group in advanced options (default is no tag). Additions can be made as required; see Tag Product Documentation for detailed information about tags. 5.Click OK to complete the creation of the security group. If the “Customized” template was selected when creating a new security group, after completion, you can click Set Rules Immediately to proceed with adding security group rules. Add Security Group Rules Overview Security groups are used to manage whether traffic from the public network or the internal network is allowed. For security reasons, most inbound directions of security groups adopt a deny access policy. If you chose the ”Open all ports” template or the ”Open ports 22, 80, 443, 3389 and ICMP protocol” template when creating a security group, the system will automatically add security group rules for certain communication ports based on the selected template type. For more details, see Overview of Security Groups. This article guides you through adding security group rules to permit or prohibit access by cloud server instances within the security group to the public network or private network. Precautions • Security group rules support IPv4 security group rules and IPv6 security group rules. • The one-click open already includes IPv4 security group rules and IPv6 security group rules. Prerequisites • You have created a security group. For specific operations, see Create a Security Group. • You already know which public network or internal network accesses need to be allowed or prohibited for the cloud server instance. For more related application cases about setting security group rules, see Security Group Application Cases. Operating Steps 1.Log in to the Cloud Server Console. 2.In the left navigation bar, click Security Groups to enter the Security Group Management page. 3.On the Security Group Management page, select the region and find the security group whose rule needs to be set. 4.In the row of the security group where the rule needs to be set, click Modify Rule under the Operation column. 5.On the Security Group Rules page, click ”Inbound Rules” and choose any of the following methods according to actual needs to complete the operation. o Method One: One-click Open, suitable for scenarios where no ICMP protocol rules need to be set and operations can be completed via ports 22, 3389, ICMP, 80, 443, 20, and 21. o Method Two: Add Rules, suitable for scenarios requiring multiple communication protocols, such as the ICMP protocol. Instructions The following steps use

Method Two: Add Rules as an example. 6.In the pop–up ”Add Inbound Rule” window, configure the rule. Main parameters for adding rules are as follows: o Type: By default, select ”Custom”. You can also choose other system rule templates, such as the ”Windows Login” template, ”Linux Login” template, ”Ping” template, ”HTTP(80)” template, and ”HTTPS(443)” template. o Source: The source of traffic (for inbound rules) or destination (for outbound rules). Please specify one of the following options:

Specified Source/Destination	Description
Single IPv4 address or IPv4 address range	CIDR notation (such as 203.0.113.0, 203.0.113.0/24, or 0.0.0.0/0, where 0.0.0.0/0 represents matching all IPv4 addresses).
Single IPv6 address or IPv6 address range	CIDR notation (such as FF05::B5, FF05:B5::/60, ::/0, or 0::0/0, where ::/0 or 0::0/0 represents matching all IPv6 addresses).
Referenced Security Group ID, you can reference the ID of the following security groups:• Current Security Group• Other Security Groups	• Current Security Group refers to the associated security group ID of the cloud server.• Other Security Groups refer to another security group ID under the same project in the same region.• Referring to a security group ID as an advanced feature, you may choose to use it. The rules of the referenced security group will not be added to the current security group.• If you input a security group ID in the source/target during configuration of security group rules, it means only using the internal IP addresses of the cloud servers and elastic NICs bound to this security group ID as the source/target, excluding external IP addresses.
Reference IP address objects or IP address groups in Parameter Templates	–

o Protocol Port: Fill in the protocol type and port range. Protocol types supported are TCP, UDP, ICMP, ICMPv6, and GRE. You can also reference protocol ports or protocol port groups in Parameter Templates. Supported formats for protocol ports are as follows: • Single port, e.g., TCP:80. • Multiple discrete ports, e.g., TCP: 80,443. • Consecutive ports, e.g., TCP:3306–20000. • All ports, e.g., TCP: ALL. o Policy: Default selection is ”Allow”. • Allow: Permit corresponding access requests for this port. • Deny: Directly discard packets without returning any response information. o Remarks: Customized, briefly describe the rule for easier management later. 7.Click Complete to finish adding the inbound rule for the security group. 8.On the Security Group Rules page, click ”Outbound Rules,” then refer to Steps 5 through 7 to complete the addition of outbound rules for the security group. Associate Instances with Security Groups Overview Security groups are used to set network access controls for single or multiple cloud server instances,

serving as important means of network security isolation. Based on business requirements, you can associate cloud server instances with one or multiple security groups. Below, we will guide you on how to associate cloud server instances with security groups on the console. Prerequisites Cloud server instances have been created. Directions 1.Log in to the Cloud Server Console. 2.In the left navigation bar, click Security Groups to enter the Security Group Management page. 3.On the Security Group Management page, select the region and locate the security group that needs rule settings. 4.In the row of the security group requiring rule settings, click Manage Instances under the Operation column to enter the Associated Instance page. 5.On the Associated Instance page, click Add Association. 6.In the pop–up ”Add Instance Association” window, select instances to be bound with the security group and click Confirm. Note After multiple security groups are bound to an instance, they will execute in order based on their binding sequence as priority order. If you need to adjust the priority of security groups, please refer to Adjust Security Group Priority. Follow–up Operations • If you want to view all security groups created under a certain region, you can query the list of security groups. For specific operations, see View Security Groups. • If you do not want your cloud server instances to belong to certain security groups, you can remove the cloud server instances from the security groups. For specific operations, see Remove from Security Group. • If your business no longer requires one or more security groups, you can delete the security groups. After deletion, all security group rules within the security group will also be deleted simultaneously. For specific operations, see Delete Security Group. Security Group Management Viewing Security Groups Overview If you want to view all the security groups created under a specific region, you can follow these steps to view the list of security groups. Steps View All Security Groups 1.Log in to the cloud server console. 2.In the left navigation bar, click on Security Groups to enter the security group management page. 3.On the security group management page, select the region to view all security groups under that region. View Specific Security Group You can also use the search function on the security group management page to view the security group you need. 1.Log in to the cloud server console. 2.In the left navigation bar, click on Security Groups to enter the security group management page. 3.On the security group management page, select the region. 4.At the top right corner of the security group list under this region, click on the search box and choose any of the following methods to query the security group you need to view. □Select Security Group ID, enter the security group ID, press Enter, and you will find the security group corresponding to this security group ID. □Select Security Group Name, enter the security group name, press Enter, and you will find the security group corresponding to this security group name. □Select Tag, enter the tag name, press Enter, and you will find all security groups under this tag. Remove from Security Group Overview You can remove instances from a security group based on business needs. Prerequisites The instance of the cloud server has already joined two or more security groups. Steps 1.Log in to the cloud server console. 2.In the left navigation bar, click on Security Groups to enter the security group management page. 3.On the security group management page, select the region and locate the security group from which you need to remove the instance. 4.Click Manage Instances in the Operations column of the row for the security group where you need to remove the instance, entering the associated instance page. 5.On the associated instance page, select the

instance to be removed and click Remove from Security Group. 6.Click OK in the pop–up dialog box. Clone Security Group Overview You may need to clone a security group when you meet the following scenarios: □Suppose you have already created a security group sg–A in Region A, and now you need to apply exactly the same rules to instances in Region B; you can directly clone sg–A to Region B without having to create a security group from scratch in Region B. □If your business requires executing a new security group rule, you can clone the original security group as a backup. Notes □Cloning a security group defaults to cloning only the inbound/outbound rules of this security group, but does not clone the instances associated with this security group. □Cloning a security group supports cross–project and cross–region operations. Steps 1.Log in to the cloud server console. 2.In the left navigation bar, click on Security Groups to enter the security group management page. 3.On the security group management page, select the region and locate the security group you need to clone. 4.In the row of the security group you need to clone, click More > Clone in the Actions column. 5.In the ”Clone Security Group” window that pops up, select the target project and target region for cloning, fill in the new name for the security group, and click OK. Delete Security Group Operating Scenario If your business no longer requires one or more security groups, you can delete the security group. After deletion, all rules within the security group will also be removed. Prerequisites Please confirm that the security group to be deleted is not associated with any instances. If there are associated instances, move them out of the security group first; otherwise, the operation to delete the security group cannot be performed. For specific operations, refer to Removing Instances from Security Groups. Steps for Deletion 1.Log in to the cloud server console. 2.In the left navigation bar, click ”Security Groups” to enter the security group management page. 3.On the security group management page, select the region and find the security group to be deleted. 4.In the row of the security group to be deleted, click More > Delete under the Operations column. 5.Click OK in the pop–up prompt box. Adjusting Security Group Priority Operating Scenario A cloud server instance can bind to one or multiple security groups. When a cloud server instance binds to multiple security groups, these security groups execute sequentially according to their priority order (such as 1, 2), and you can adjust the priority of the security groups based on the following steps. Prerequisites The cloud server instance has joined two or more security groups. Steps for Adjustment 1.Log in to the cloud server console. 2.On the Instance Management page, click the ID of the cloud server instance to enter its details page. 3.Select the Security Groups tab to enter the security group management page. 4.In the ”Bound Security Groups” module, click

Sort.

5.Click the icon below and drag it up and down to adjust the priority of the security group—the higher the position, the higher the priority of the security group.

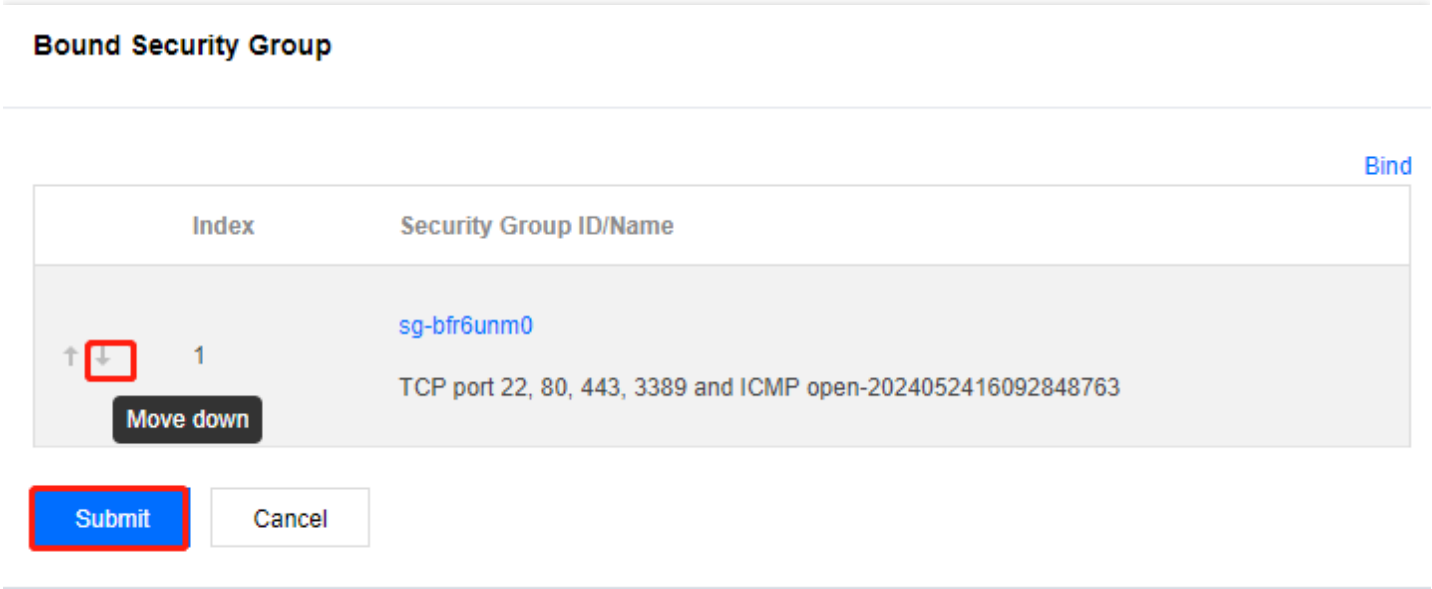

6.After completing the adjustment, click Save. Manage Security Group Rules View Security Group Rules Operating Scenario After adding security group rules, you can view the detailed information about the security group rules on the console. Prerequisites A security group has been created, and security group rules have been added to this security group. For instructions on creating a security group and adding security group rules, see Create Security Group and Add Security Group Rules. Steps for Viewing 1.Log in to the Cloud Server Console. 2.In the left navigation bar, click Security Groups to enter the security group management page. 3.On the security group management page, select the region and locate the security

group whose rules you wish to view. 4.Click the ID or name of the security group whose rules you wish to view to enter the security group rule page. 5.On the security group rule page, click the inbound or outbound rule tabs to view the corresponding security group rules. Modify Security Group Rules Operating Scenario Incorrectly configured security group rules can pose significant security risks, such as permitting uncontrolled access to specific ports. By modifying unreasonable security group rules within the security group, you can ensure the network security of cloud server instances. This guide instructs you on how to modify security group rules. Prerequisites A security group has been established, and security group rules have been added to this security group. Refer to Create Security Group and Add Security Group Rules for guidance on setting up a security group and adding security group rules. Steps for Modification 1.Log in to the Cloud Server Console. 2.In the left navigation bar, click Security Groups to enter the security group management page. 3.On the security group management page, choose the region and identify the security group requiring rule modifications. 4.In the row of the security group needing rule changes, click Modify Rule under the Actions column to open the security group rule page. 5.On the security group rule page, select the inbound or outbound rule tabs based on the direction of the security group rule to be modified. Deleting Security Group Rule Overview If you no longer need a security group rule, you can delete it. Prerequisites ●A security group has been created and security group rules have been added to it. For information on how to create a security group and add security group rules, see Create a Security Group and Add Security Group Rules. ●It is confirmed that the CVM instance does not need to permit/forbid public network access or private network access. Directions 1.Log in to the CVM console. 2.Click Security Group in the left sidebar to enter the security group management page. 3.Select Region on the security group management page and find the security groups whose rules need to be deleted. 4.In the row of the security group whose rules need to be deleted, click Modify Rule in the operation column to enter the Security Group Rules page. 5.On the Security Group Rules page, depending on the direction of the rules to be deleted (Inbound/Outbound), click the Inbound Rules or Outbound Rules tab. 6.Find the security group rule that needs to be deleted and click Delete in the operation column. 7.In the pop–up dialog Box, click OK. Exporting Security Group Rule Overview Security group rules can be exported. You can export the security group rules of a security groups for local backup. Directions 1.Log in to the CVM console. 2.Click Security Group in the left sidebar to enter the security group management page. 3.Select Region on the security group management page and find the security groups whose rules need to be exported. 4.Click the security group ID/name of the rule to be exported to enter the Security Group Rules page. 5.On the Security Group Rules page, depending on the direction of the rules to be exported (Inbound/Outbound), click the Inbound Rules or Outbound Rules tab. 6.Under the Inbound/Outbound Rules tab, click the Download Inbound Rules or Outbound Rules icon at the top right to download and save the security group rule file to your local device. Importing Security Group Rule Overview Security group rules can be imported. You can import the exported security group rule file into a security group to quickly create or recover security group rules. Directions 1.Log in to the CVM console. 2.Click Security Group in the left sidebar to enter the security group management page. 3.Select Region on the security group management page and find the security groups whose rules need to be imported. 4.Click the security group ID/name of

the rule to be imported to enter the Security Group Rules page. 5.On the Security Group Rules page, depending on the direction of the rules to be imported (Inbound/Outbound), click the Inbound Rules or Outbound Rules tab. 6.Click Import Rule under the Inbound/Outbound Rules tab. 7.In the pop–up Batch Import – Inbound/Outbound Rules window, select the edited template file for inbound/outbound rules, and click Start Importing. Note: –If there are security group rules under the security groups that need to be imported, it is recommended that you export the existing rules first. Otherwise, when importing new rules, the original rules will be overwritten. –If there is no security group rule under the security group that needs to import rules, it is recommended that you download the template first and then import the file after editing the template file. Common Server Ports

Port	Service	Note
21	FTP	Port opened by the FTP server. It is for uploading and downloading.
22	SSH	Port 22 is the SSH port used to remotely connect Linux system servers via command–line mode.
25	SMTP	Port opened by the FTP server. It is for sending mail.
80	HTTP(80)	It is used to provide external access for website services such as IIS, Apache, and Nginx.
110	POP3	Port 110 is open for POP3 (mail protocol 3) services.
137,138,139	NETBIOS Protocol	Among them, 137 and 138 are UDP ports, which are used when transferring files through the network neighborhood.Port 139: Connections coming through this port attempt to get NetBIOS/SMB services. This protocol is used for Windows file and printer sharing and SAMBA.
143	IMAP	Port 143 is mainly used for Internet Message Access Protocol v2 (IMAP). Like POP3, it is a protocol for receiving emails.
443	HTTPS	A web browsing port and another type of HTTP that provides encryption and transmission through secure ports.
1433	SQL Server	Port 1433 is the default port of the SQL Server, which uses two ports: TCP–1433 and UDP–1434. Port 1433 is used for SQL Server to provide external services, whereas port 1434 is used to respond to the requester regarding which TCP/IP port is being used by SQL Server.
3306	MySQL	Port 3306, the default port of MySQL database, is used for MySQL to provide external services.

Port	Service	Note
3389	Windows Server Remote Desktop Services	Port 3389 is the service port for Windows 2000 (2003) Server remote desktop, which allows connection to a remote server using the Remote Desktop connection tool.
8080	Proxy Port	Similar to port 80, port 8080 is used for the WWW proxy service for web browsing. The port number extension ":8080" is often appended to the URL when users visit a website or use a proxy server. In addition, after the Apache Tomcat web server is installed, the default service port is port 8080.

Security Group API Overview

Interface Name	Interface Function
CreateSecurityGroup	Create a security group
CreateSecurityGroupPolicies	Add rules to a security group
DeleteSecurityGroup	Delete a security group
DeleteSecurityGroupPolicies	Delete security group rules
DescribeSecurityGroupAssociationStatistics	Query statistics on instances associated with a security group
DescribeSecurityGroupPolicies	Query security group rules
DescribeSecurityGroups	View security groups
ModifySecurityGroupAttribute	Modify attributes of a security group
ModifySecurityGroupPolicies	Modify outbound and inbound rules of a security group
ReplaceSecurityGroupPolicy	Replace a single route rule of a security group

Manage SSH Key To ensure the security and reliability of the instance, the cloud platform provides two encrypted login methods: password login and SSH key pair login. This document describes common

operations on SSH key pairs. Creating SSH Keys 1.Log in to the CVM Console. 2.Click SSH Key in the left

sidebar. 3.Click Create Key.

○Creation Method: If you select "Create a new key pair," then enter the key name. If you select "Import existing public key," then enter the key name and the existing public key information. Note: You must use a public key without a password; otherwise, you will not be able to successfully log in to the instance through the console. ○Key Name: Customized name. ○Label (optional): You may add labels to keys as needed for resource classification, search, and aggregation. For more information, see Labels. Click OK to complete the creation. Note: After clicking OK, the private key will be automatically downloaded. The cloud platform does not retain your private key information, so please keep your private key properly. Key binding instance 1.Log in to the CVM Console. Click SSH Key in the left sidebar. 2.Select the SSH key and click Bind CVM.

3.Select a region, check the CVM that you want to associate, and click Confirm. Key unbinding instance 1.Log in to the CVM Console. Click SSH Key in the left sidebar. 2.Select the SSH key and click Unbind

CVM.

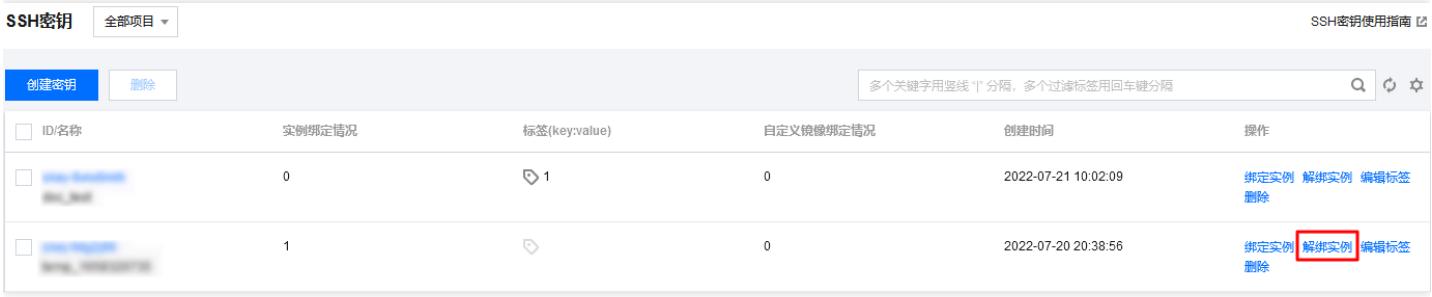

3.Select a region, check the CVM that you want to unbind, and click Confirm. Change the SSH key name/description 1.Log in to the CVM Console. Click SSH Key in the left sidebar. 2.Select the key to be modified in the key list and click Modify above.

3.Enter the new name and description in the pop-up dialog box and click Confirm. Deleting SSH Keys Note: If the SSH key has been associated with a CVM or a custom image, it cannot be deleted. 1.Log in to the CVM Console. Click SSH Keys in the sidebar. 2.On the "SSH Keys" page, you may delete keys as needed: ○Select the Delete option next to the SSH key you wish to remove, as shown in the image below:

○In the pop-up window for deleting the key, click OK.

Tag

Tag

Last Updated At: 2025-08-14 09:25:05

Managing Instances Using Tags Overview Tag is a key-value pair to identify cloud resources. Tags can help you conveniently classify and manage CVM resources from various dimensions (such as business, purpose, owner, etc.). It should be noted that the tags you set are only used for your management of server resources. Use Limits You need to pay attention to the following use limits of tags: ●Quantity limit: Each cloud resource can have up to 50 tags. ●Tag Key Limits: -qcloud, project is reserved tag keys by the system and cannot be created. -Only digits, letters, and +=.@- are allowed, and the length of the tag key cannot exceed 255 characters. ●Tag value limit: It can only be null character strings or digits, letters, and +=.@-, and the maximum length of the tag value is 127 characters. Directions and Cases Case Description Example: A company purchased 6 CVM instances, and the information of the use department, business scope, and owner of these 6 instances is as follows.

Instance instance-id	Department	Business Scope	Owner
ins-abcdef1	E-commerce	Marketing Campaigns	Tom
ins-abcdef2	E-commerce	Marketing Campaigns	Harry
ins-abcdef3	Game	Game A	Jane
ins-abcdef4	Game	Game B	Harry
ins-abcdef5	Entertainment	Post-production	Harry
ins-abcdef6	Entertainment	Post-production	Tom

Taking ins-abcdef1 as an example, we can add the following three sets of tags to the instance:

tag key	tag value
dept	ecommerce
business	mkt
owner	zhangsan

Similarly, you can also set appropriate tags for other instances based on their use department, business scope, and owner. Setting Tags in CVM Console After designing the tag keys and values as detailed

above, you can log in to the Cloud Platform CVM console to set tags. 1.Log in to the CVM Console. 2.In the instance management console, select the instance you want to edit tags for, click [More]>[Cloud Host Settings]>[Edit Tag].

3.In the pop–up window titled [You have selected 1 cloud resource], set the tags. As shown below: For example, add three sets of tags to the instance ins–abcdef1.

4.Click OK. Filtering Instances by Tags When you want to filter instances of a certain type of tags, you can follow the steps below. 1.In the search box, select Tags. 2.Enter the tag key and tag value after Tag: and click to search. For example, if you want to filter out CVM resources whose owner is Tom, you can enter tag:owner:Tom. Editing Tags Overview

Use Limits When editing labels, you need to be aware of the following limits: ●Quantity limit: Each cloud resource can have up to 50 tags. ●Tag Key Limits: –qcloud, project is reserved tag keys by the system and cannot be created. –Only digits, letters, and +=.@– are allowed, and the length of the tag key cannot exceed 255 characters. ●ag value limit: It can only be null character strings or digits, letters, and +=.@–, and the maximum length of the tag value is 127 characters. Prerequisites You have logged in to the CVM console. Directions Editing Tags for Individual Instances 1.On the Cloud Host list page, select the instances whose tags you want to edit, click More > Cloud Host Settings > Edit Tag. 2.In the pop–up You have selected 1 resource window, you can add, modify or delete tags based on actual needs.

Example of Cloud Access Management

Last Updated At: 2025-08-11 14:49:05

Overview You can use Cloud Access Management (CAM) policies to allow users to view and use specific resources in the CVM console. This document provides examples of permissions for checking and using specific resources, and policies that guide users on how to use specific parts of the console. Operation Example Full Read-Write Policy of CVM If you want to allow a user to have the permission to create and manage CVM instances, you can use the policy named CVMFullAccess for the user. This policy is achieved by allowing users to have operation permissions on all resources in CVM, VPC, CLB and MONITOR. Read-Only Policy of CVM If you want to allow a user to have the permission to query CVM instances only but not the permission to create, delete, or power on/off instances, you can use the policy named CVMInnerReadOnlyAccess for the user. This policy is achieved by allowing users to have the following permissions for all operations that start with the word Describe and all operations that start with the word Inquiry in CVM. Read-Only Policy for CVM-Related Resources If you want to allow a user to have the permission to query CVM instances and related resources (VPC, CLB) only, but not the permission to create, delete, power on/off instances, you can use the policy named CVMReadOnlyAccess for the user. This policy is achieved by allowing users to have operation permissions for the following operations: ●All operations in CVM that start with the word Describe and all operations that start with the word Inquiry. ●All operations in VPC that start with the word Describe, all operations that start with the word Inquiry, and all operations that start with the word Get. ●All operations in the CLB that start that start with the word Describe. ●All operations in Monitor that start with the word Describe and all operations that start with the word Get. The directions are as follows: Referring to the authorization management, grant the preset policy QcloudCVMReadOnlyAccess to the user. Policies Related to Elastic Cloud Storage If you want the user to be able to view the Cloud Block Storage information in the CVM console and have permissions to create and use Cloud Block Storage, you can first add the following operations to your policy and then associate the policy with the user. ●CreateCbsStorages: Create Cloud Block Storage. ●AttachCbsStorages: Mount the specified elastic cloud storage to the specified Cloud Virtual Machine. ●DetachCbsStorages: Unmount the specified elastic cloud storage. ●ModifyCbsStorageAttributes: Modify the name or project ID of the specified Cloud Block Storage. ●DescribeCbsStorages: Query the detailed information of the Cloud Block Storage. ●DescribeInstancesCbsNum: Query the number of elastic cloud storage mounted on the Cloud Virtual Machine and the total number of elastic cloud storage that can be mounted. ●RenewCbsStorage: Renew the specified elastic cloud storage. ●ResizeCbsStorage: Scale out the specified elastic cloud storage. The directions are as follows: 1.According to the strategy, create a custom policy that can view Cloud Block Storage information in the CVM console and has other permissions such as creating Cloud Block Storage and using Cloud Block Storage. The policy content can be set by referring to the following policy syntax: { ”version”: ”2.0”, ”statement”: [{ ”effect”: ”allow”, ”action”: [”name/cvm:CreateCbsStorages”,

”name/cvm:AttachCbsStorages”, ”name/cvm:DetachCbsStorages”, ”name/cvm:ModifyCbsStorageAttributes”, ”name/cvm:DescribeCbsStorages”], ”resource”: [”qcs::cvm::uin/1410643447:*”] }] }

2.Find the created policy, and in the Operation column of the policy row, click Associate User/Group/Role. 3.In the pop-up Associate User/User Group/Role window, select the user/user group you want to authorize and click OK. Policies Related to Security Group If you want the user to be able to view and use the security groups in the CVM console, you can add the following operations to your policy and then associate the policy with the user.

●DeleteSecurityGroup: Delete the security group. ●ModifySecurityGroupPolicys: Replace all policies of the security group. ●ModifySingleSecurityGroupPolicy: Modify a single policy of the security group. ●CreateSecurityGroupPolicy: Create the security group policy. ●DeleteSecurityGroupPolicy: Delete the security group policy. ●ModifySecurityGroupAttributes: Modify the security group attributes. The directions are as follows: 1.According to the strategy, create a custom policy that allows users to have other permissions such as creating, deleting, and modifying security groups in the CVM console. The policy content can be set by referring to the following policy syntax: { ”version”: ”2.0”, ”statement”: [{ ”action”: [”name/cvm:ModifySecurityGroupPolicys”, ”name/cvm:ModifySingleSecurityGroupPolicy”, ”name/cvm:CreateSecurityGroupPolicy”, ”name/cvm:DeleteSecurityGroupPolicy”], ”resource”: ”*”, ”effect”: ”allow” }] } *2.Find the created policy, and in the Operation column of the policy row, click Associate User/Group/Role. 3.In the pop-up Associate User/User Group/Role window, select the user/user group you want to authorize and click OK. Policies Related to Elastic IP Address If you want the user to be able to view and use elastic IP addresses in the CVM console, you can add the following operations to your policy and then associate the policy with the user. ●AllocateAddresses: Allocate the address to VPC or CVM. ●AssociateAddress: Associate the elastic IP address with an instance or a network API. ●DescribeAddresses: View the elastic IP address in the CVM console. ●DisassociateAddress: Disassociate the elastic IP address with an instance or a network API. ●ModifyAddressAttribute: Modify the attributes of the elastic IP address. ●ReleaseAddresses: Release the elastic IP address. The directions are as follows: 1.According to the strategy, create a custom policy. This policy allows users to view elastic IP addresses in the CVM console and assign and associate them with instances, but does not allow users to modify the attributes of elastic IP addresses, disassociate elastic IP addresses, or release elastic IP addresses. The policy content can be set by referring to the following policy syntax: { ”version”: ”2.0”, ”statement”: [{ ”action”: [”name/cvm:DescribeAddresses”, ”name/cvm:AllocateAddresses”, ”name/cvm:AssociateAddress”], ”resource”: ”*”, ”effect”: ”allow” }] }*
2.Find the created policy, and in the Operation column of the policy row, click Associate User/Group/Role. 3.In the pop-up Associate User/User Group/Role window, select the user/user group you want to authorize and click OK. Authorizing Users to Have Specific CVM Operation Permissions To authorize a user specific CVM operation permissions, you can associate the following policy with the user. The directions are as follows: 1.According to the strategy, create a custom policy. This policy allows users to

have operation permissions on the CVM instance with ID ins–1 and region Guangzhou. The policy content can be set according to the following policy syntax: { ”version”: ”2.0”, ”statement”: [{ ”action”: ”cvm:*”, ”resource”: ”qcs::cvm:ap–cityA::instance/ins–1”, ”effect”: ”allow” }] } 2.Find the created policy, and in the Operation column of the policy row, click Associate User/Group/Role. 3.In the pop–up Associate User/User Group/Role window, select the user/user group you want to authorize and click OK. Authorizing Users to Have CVM Operation Permissions for Specific Regions To authorize a user’s CVM operation permissions for a specific region, you can associate the following policy with the user. The directions are as follows: 1.According to the strategy, create a custom policy. This policy allows users to have operation permissions on CVM machines in the Guangzhou region. The policy content can be set according to the following policy syntax: { ”version”: ”2.0”, ”statement”: [{ ”action”: ”cvm:*”, ”resource”: ”qcs::cvm:ap–cityA::*”, ”effect”: ”allow” }] } 2.Find the created policy, and in the Operation column of the policy row, click Associate User/Group/Role. 3.In the pop–up Associate User/User Group/Role window, select the user/user group you want to authorize and click OK.

Authorizing the Sub–Account to Have All the Permissions of CVM but Not Including Payment Permissions Assume that there is a sub–account (Developer) under the enterprise account (CompanyExample, ownerUin is 12345678). The sub–account needs to have all management permissions (such as creation, management, and all other operations) for the CVM service of the enterprise account, but does not include payment permissions (one can place orders but cannot pay). We can achieve this through the following two solutions: ●Solution A The enterprise account CompanyExample directly authorizes the preset policy QcloudCVMFullAccess to the sub–account Developer. ●Solution B Create a custom policy using the following policy syntax. { ”version”: ”2.0”, ”statement”:[{ ”effect”: ”allow”, ”action”: ”cvm:*”, ”resource”: ”*” }] } uthorize the policy to the sub–account. Authorizing the Sub–account to Have the Operation Permission of Project Management Assume that there is a sub–account (Developer) under the enterprise account (CompanyExample, ownerUin is 12345678). The sub–account needs to be authorized to manage resources in the console based on the project. The directions are as follows: 1.Create a custom policy for project management based on business permissions. 2.Authorize the created custom policy to the sub–account. If a sub–account is prompted for no permission when managing a project, such as viewing snapshots, images, VPCs, Elastic IPs, you can authorize the preset policies QcloudCVMAccessForNullProject, QcloudCVMOrderAccess, and QcloudCVMLaunchToVPC to the sub–account.

Overview

Last Updated At: 2025-08-11 14:58:27

This document introduces common scenarios and related operations when using Cloud Virtual Machine instances and Cloud Virtual Machine related products for your reference. Purchasing and Using a CVM If you purchase and use Cloud Virtual Machine for the first time, it is recommending that you learn about, purchase, and use it in the following order. 1. To understand the concept of Cloud Virtual Machine, see CVM Overview. 2. Cloud Virtual Machine selection and purchase. If you are an individual user and use it for the first time, it is recommended that you use the quick configuration method to Customizing Windows CVM Configurations or Customizing Linux CVM Configuration. 3. After the purchase is completed, log in to the Cloud Virtual Machine: Depending on the type of Cloud Virtual Machine you purchased, you can choose Log in to the Windows Instance or Log in to the Linux Instance. 4. (Optional) You can use the purchased Cloud Virtual Machine to build a personal website, forum, or storage files. Build a personal WordPress site using an image Set up a Discuz! forum using an image Adjusting CVM Configurations After you purchase a Cloud Virtual Machine, you may need to adjust the hard disk, network, and other configurations of the Cloud Virtual Machine due to changes in requirements. You can see the following documents to complete the operation. Adjusting instance configuration Adjusting network configuration Adjusting project configuration Reinstalling system Resetting Password and Key If you forget your Cloud Virtual Machine password or lose your key, you can see the following document to reset your password or key. ● Resetting instance password ● Managing SSH keys Creating, Importing or Deleting a Custom Image Image provides all the information required to launch a Cloud Virtual Machine instance. Generally speaking, the image is the installation disk of the Cloud Virtual Machine. Currently, Cloud Platform provides three types of images: public images, custom images, and shared images. The following describes the common operations supported by the image. ● Creating custom image ● Deleting custom image ● Importing image ● Copying image Troubleshooting If you are unable to log in to a Cloud Virtual Machine or the Cloud Virtual Machine is operating slowly, you can see the following documents for troubleshooting. CVM Login Failures Cloud Virtual Machine network latency and packet loss

Use Limits

Last Updated At: 2025-08-14 09:26:49

Usage limitations of CVM instances: □Currently does not support installation of virtualization software and re-virtualization (such as installing and using VMware or Hyper-V). □Does not currently support sound card applications, direct loading of external hardware devices (such as USB drives, external hard disks, bank USB tokens, etc.). □The public network gateway currently supports Linux systems only. Limits on Images ●Public images: No use limits. ●Custom images: Each region supports up to 50 custom images. ●Shared images: Each custom image can be shared with up to 50 Cloud Platform users, and can only be shared with accounts in the same region. ●For more details, see Image Type Limits. Limits on Disk

Limit Type	Limit Description
Elastic Cloud Disk Capability	Data disks applied for together with cloud servers are all elastic cloud disks, supporting detachment from and reattachment to cloud servers. This feature is supported in all availability zones.
Cloud Disk Performance Limits	I/O performance limits apply simultaneously. For example, a 1TB SSD cloud disk can achieve a maximum random IOPS of 26,000, meaning both read IOPS and write IOPS can reach this value. However, due to multiple performance constraints, in this case, using I/O with a block size of 4KB/8KB can reach the maximum IOPS value, but using I/O with a block size of 16KB cannot reach the maximum IOPS value (throughput has already reached the limit of 260MB/s).
Number of Elastic Cloud Disks Mountable to a Single Cloud Server	up to 20
Snapshot Quota per Region	64 + number of cloud disks in region * 64 (snapshots)
Restrictions on Cloud Servers where Cloud Disks can be Mounted	Cloud servers and cloud disks must be in the same availability zone.

Limit Type	Limit Description
Rollback Limitations of Snapshots	Snapshot data can only be rolled back onto the source cloud disk used to create the snapshot.
Type Restriction of Cloud Disks Created from Snapshots	Only snapshots of data disks can be used to create new elastic cloud disks.
Size Restrictions of Cloud Disks Created from Snapshots	The capacity of new cloud disks created from snapshots must be greater than or equal to the capacity of the source cloud disk.
Recovery of Overdue Cloud Disks	If prepaid elastic cloud disks expire and are not renewed within seven days after expiration, they will be recovered into the recycle bin. After entering the recycle bin, the attachment relationship between the cloud disk and the cloud server is not actively terminated. For details on the recovery mechanism, refer to the overdue payment instructions. Currently, when attaching prepaid elastic cloud disks to prepaid cloud servers, you can choose the following renewal methods based on actual needs: • Align with the expiration date of the cloud server. • Automatically renew monthly after the cloud disk expires. • Directly attach without handling renewal.

Limits on Security Group □Security groups are divided by region. A Cloud Virtual Machine (CVM) can be bind only to security groups in the same region. □The security group are applicable to any Cloud Virtual Machine instance in the network environment. □Each user can set up to 50 security groups under each project in each region. □A maximum of 100 inbound or outbound access policies can be set for each security group. □One Cloud Virtual Machine can join multiple security groups, and one security group can be associated with multiple Cloud Virtual Machines at the same time. □The security group bound to the Cloud Virtual Machine in the Basic Network cannot be filtered data packets from (or to) the relational database (CDB) and elastic cache (Redis and Memcached) on Cloud Platform. If you need to filter traffic for such instances, you can use iptables to do so. □Related quota limits are as follows:

Feature Description	Quantities
Security Group	50/region
Access Policy	100/inbound, 100/outbound
Number of cloud server instances associated with a single security group	2000
Number of security groups each cloud server instance can be associated with	5
Number of security group IDs that can be referenced by each security group	10

Restrictions related to VPC

Resource	Limit (number)
Number of private networks within each region per account	20
Number of subnets within each private network	100
Number of basic network hosts supported by association within each private network	100
Number of route tables within each private network	10
Number of route tables associated with each subnet	1
Number of routing policies for each route table	50
Default quota number of HAVIP for each private network	10

Instance
Reinstalling System

Last Updated At: 2025-08-14 09:25:05

Overview Reinstalling the system can restore the instance to its initial state after it is started. It is an important recovery method when the instance encounters a system failure. This document guides you on how to reinstall the operating system. The Cloud Virtual Machine provides the following two types of reinstallation, and both types of reinstallation can be performed on Cloud Virtual Machines in any region. Notes ●Reinstallation preparation: The contents of the system disk are lost after reinstallation, so you need to back up important information on the system disk before reinstallation. If you need to preserve system running data, it is recommended that you create a custom image before you reinstall the system and select the image for reinstallation. ●Image selection suggestions: It is recommended to use the image provided by Cloud Platform or a custom image for reinstallation. It is not recommended to use images from unknown sources or other sources. Do not perform other operations while reinstalling the system disk. ●Instance physical properties: The public network IP of the instance will not be changed. ●Instance specification restrictions: If you need to reinstall your instance using Windows 2016 and 2019 related image versions, the instance memory must be greater than 2GB. Directions You can reinstall the operating system in the following ways: 1.Log in to the CVM Console. 2.In the instance row where the system needs to be reinstalled, click Action > Reinstall System. 3.In the popped-up ”Reinstall System” window, after reading the ”System Reinstallation Instructions,” click Next. 4.select to use the current machine image or another image, adjust the disk size, enter the password, and click Start Rebuilding.

Terminating Instance

Last Updated At: 2025-08-14 09:25:05

Overview When you no longer need an instance, you can destroy it and the services running on the destroyed instance will be terminated. Directions 1.Log in to the CVM Console. 2.Select different instance destruction operations based on actual needs. -Destroy Single Instance: Find the instance to be destroyed in the list, then on the right side click Action > CVM Status > Terminate. -Bulk Destroy Instances: Check all the instances that need to be destroyed, click on the top More > Terminate. Instances that cannot be destroyed will display the reason. 3.Review the resources that will be destroyed and retained, after selecting Next, click Confirm.

Modify Instance Name

Last Updated At: 2025-08-14 09:26:40

Overview To facilitate user management of cloud server instances on the cloud server console, allowing quick identification of each cloud server instance by name, the cloud platform supports naming each instance and allows changes at any time that take effect immediately. Operation Steps Modifying the Name of a Single Instance 1.Log in to the Cloud Server Console. 2.In the instance list, select the row corresponding to the cloud server whose instance name needs modification, then click Operations > Cloud Host Settings > Rename on the right side. 3.In the pop-up "Rename" window, enter the new instance name, and click OK to confirm.

Resetting Instance Password

Last Updated At: 2025-08-14 09:26:40

Overview If you have forgotten your password, you can reset the instance's log-in password in the console. This document only describes how to change the instance log-in password in the Cloud Virtual Machine management console. 1.Log in to the cloud server console. 2.On the instance management page, operate according to the actual view mode being used: Choose the More > Password/Key > Reset Password option next to the row of the cloud server whose password needs resetting. 3.In the "Set Password" step, select the type "Username," enter the username that requires password reset, along with the corresponding "New Password" and "Confirm Password," then click Next.

Managing Instance IP Address

Last Updated At: 2025-08-14 09:26:40

Getting the Private IP Address and Set DNS This document describes how to obtain the private IP address of the instance and how to set private network DNS. Obtain the Private IP Address of the Instance 1.Log in to the CVM Console. 2.On the instance management page, select the instance for which you need to view the Private IP, move the mouse to the Primary IPv4 Address column, a will appear. Click to copy the IP address. Set Private Network DNS When network resolution errors occur, you can manually set private network DNS based on the type of Cloud Virtual Machine operating system. 1.Log in to the Linux Cloud Virtual Machine. 2.Execute the following command to open /etc/sysctl.conf file.

3.Press i to switch to Edit mode, and modify DNS IP according to the corresponding region in the Private Network DNS list. 4.Press Esc, enter :wq to save the file and return. Modifying Private IP Addresses Overview You can directly modify the private IP address of a Cloud Virtual Machine (CVM) instance in the Virtual Private Cloud in the console, or you can change the private IP address of an instance by changing the subnet to which the CVM instance belongs. This document guides you to modify the private IP address of a CVM instance in the Virtual Private Cloud in the Cloud Virtual Machine console. For more information about how to change the subnet, see Change Instance Subnet. Limitation Factor ●Changing the primary IP of the primary network interface will cause the associated Cloud Virtual Machine to automatically restart. ●The secondary network interface cannot modify the primary IP. Directions 1.Log in to the CVM Console. 2.Select the region where the instance whose private IP address is to be modified belongs, and click ID/Host Name of the instance to enter the Instance Details page. 3.On the Instance Details page, select the ENI tab, click Change Primary IP. 4.In the pop-up Change Primary IP window, enter the new IP, click Confirm, and wait for the instance to restart for the changes to take effect.

Changing Instance Subnet

Last Updated At: 2025-08-14 09:26:41

Overview This document guides you to directly change the subnet to which a Cloud Virtual Machine instance in a Virtual Private Cloud belongs in the console. Limitation Factor ●Changing the subnet will cause the associated Cloud Virtual Machine to automatically restart. ●The secondary network interface cannot change the subnet. Directions 1.Log in to the CVM Console. 2.On the CVM host list page, select the region where the instance whose subnet needs to be changed belongs. 3.On the CVM host list page, find the instance whose subnet needs to be changed, and click its ID/hostname to enter the instance details page. 4.On the instance details page, select the ENI tab, and click Primary ENI. Enter the Primary Network Interface Management page. 5.On the ENI Management Page, click Change Subnet. 6.In the popped-up Change Subnet window, select the new subnet, enter the new primary IP, and click OK. After the instance is restarted, it takes effect.

Changing the Security Group

Last Updated At: 2025-08-14 09:26:41

Overview A security group is a stateful virtual firewall with packet filtering. It is used to set network access control for one or more CVMs and is an important network security isolation method provided by Cloud Platform. Creating a CVM instance necessitates assigning a security group to the instance. Cloud Platform allows users to change the security group to which the instance belongs after creating the CVM instance. Prerequisites You have logged in to the Cloud Virtual Machine console. Changing a Configured Security Group You have logged in to the Cloud Virtual Machine console. 1.In the CVM list page, select a CVM that needs to be reassigned to a new security group, then click Action > Configure Security Group. 2.In the pop-up Configure Security Group window, check the new security group names (multiple selections allowed), click Confirm to complete the security group change. Changing a Bound Security Group 1.In the CVM list page, click on the CVM instance ID/hostname that needs a security group assigned, to enter the instance detail page. 2.On the Instance Details Page, select the Security Group tab, and in the ”Bound Security Group” column, click Bind. 3.In the pop-up Configure Security Group window, select the security groups you want to bind as needed, click Confirm to complete the binding.

Searching for Instance

Last Updated At: 2025-08-14 09:26:41

Overview By default, the Cloud Virtual Machine console displays the Cloud Virtual Machine of all projects in the current region. To help users quickly search for Cloud Virtual Machines in the current region, Cloud Platform provides a Cloud Virtual Machine search feature, which can currently be filtered by resource attribute dimensions such as project, instance type, availability zone, IP, instance ID, and instance name. Directions 1.Log in to the CVM Console. 2.In the search box, enter the content you want to search for according to your actual needs and click to search. 3.Select a searchable resource dimension (such as host name, host ID, host type), and click .

Exporting Instance

Last Updated At: 2025-08-14 09:26:41

Overview You can export a list of Cloud Virtual Machine instances in a certain region in the console and customize the fields in the exported list. You can select up to 26 fields for custom export fields. The fields currently supported for export include: ID, host name, status, region, availability zone, host type, operating system, image ID, CPU (cores), memory (G), bandwidth (Mbps), primary IPv4 public network IP, primary IPv4 private network IP, primary IPv6 address, system disk type, system disk size (GB), data disk type, data disk size, network, subnet, associated vpc, project, creation time, billing mode, tag, and placement group. Directions 1.Log in to the CVM Console. 2.Select region. 3.Click in the upper right corner of the instance list.

Turn On Instance

Last Updated At: 2025-08-14 09:26:41

Overview This article introduces how to start instances in the shutdown state through the cloud server console and cloud API. Directions Turning on a single instance 1.Log in to the cloud server console. 2.Select the instance to be started, and in the right operation column, choose Operation > Cloud Host Status > Turn On. Turning on multiple instances 1.Log in to the cloud server console. 2.Check all the instances that need to be turned on, and click Turn On at the top of the list to batch turn on instances.

Shutting Down Instance

Last Updated At: 2025-08-14 09:26:41

Overview If you need to stop the instance service or need to perform a configuration that can only be modified when the instance is shut down, you can shut down the instance. Shutting down an instance is equivalent to shutting down the local computer. Notes ●You can use system commands to shut down the instance (such as the shutdown command in Windows and the shutdown command in Linux), or you can use the Cloud Platform console to shut down the instance. It is recommended to open the console during shutdown to view the shutdown process to see if there is any problem. ●After the instance is shut down, services will not be available. Therefore, before the shutdown, ensure that the Cloud Virtual Machine has suspended business requests. ●The instance is shut down normally. The status changes to Shutting Down first, and then changes to Shutdown after the shutdown is complete. If the shutdown time is too long, problems may occur. For details, see Shutdown Related to avoid forced shutdown. ●After the instance is shut down, all storage remains connected to the instance and all disk data is preserved. Data in memory will be lost. ●Shutting down an instance does not change the physical characteristics of the instance. The public network IP and private network IP of the instance remain unchanged; Elastic IP maintains the binding relationship, but due to the service interruption, you will get an error response when accessing these IPs. ●If the shutdown instance belongs to the backend server cluster of the Cloud Load Balancer instance, after shutdown, services will not be available.If a health check policy is configured, the shutdown instance can be automatically masked and no requests will be forwarded to it. If no health check policy is configured, the client may receive a 502 error response. For more information, see Health Check. Directions Shut down a single instance 1.Log in to the CVM Console. 2.Select the instance you want to shut down, and in the right operation column, click Action > CVM Status > Shut Down. Shut down multiple instances 1.Log in to the CVM Console. 2.Check all instances that need to be shut down, then at the top of the list, click Shut Down to bulk shut down the instances. Instances that cannot be shut down will display the reason.

Restarting Instances

Last Updated At: 2025-08-14 09:26:41

Restarting is a common way to maintain the Cloud Virtual Machine. Restarting an instance is equivalent to restarting the operating system of a local computer. Overview □Restart preparation: The instance can't provide services normally during the restart, so before it is restarted, make sure that the Cloud Virtual Machine has suspended business requests. □Restart operation mode: It is recommended to use the restart operation provided by the cloud platform to restart the instance instead of running the restart command in the instance (such as the restart command in Windows and the Reboot command in Linux). □Restart time: Generally speaking, the restart operation only takes a few minutes. □Instance physical characteristics: Restarting an instance does not change the instance's physical characteristics. The instance's public network IP, private network IP, and any stored data do not change. Directions You can restart the instance in the following ways: Restart a single instance 1.Log in to the CVM console. 2.Select the instance you need to restart, click Restart at the top of the list, or choose Action > CVM Status > Restart from the right-side operations column. Restart multiple instances 1.Log in to the CVM console. 2.Select all the instances you need to restart, and click Restart at the top of the list. You can then bulk restart the instances. Instances that cannot be restarted will display the reason.

Creating an instance

Last Updated At: 2025-08-14 09:26:49

Creating an instance through the application page Overview This document uses the custom configuration method as an example to guide you on how to create a Cloud Virtual Machine (CVM) instance. Prerequisites Before creating a CVM instance, you need to complete the following tasks: Create a cloud platform account. To create a CVM instance with a network type of Virtual Private Cloud, you need to create a Virtual Private Cloud in the target region and a subnet in the target availability zone within the Virtual Private Cloud. If you do not want to use the system's automatically created default security group, you need to create a security group in the target region and add security group rules that meet your business needs. If you need to bind an SSH key pair when creating a Linux instance, you need to create an SSH key in the target project. If you need to create a CVM instance with a custom image, you need to create a custom image or import an image. Directions 1.Log in to the cloud platform, select Cloud Services> Cloud Computing and Networking > Cloud Virtual Machine > CDH, click Create to enter the CVM purchase page. –Custom configuration: Suitable for use in specific scenarios, making it convenient for users to purchase CVM instances that meet their specific needs. 2.Follow the on–page prompts to configure the following information:

Category	Required/Optional	Configuration Description
Region/Availability Zone	Required	Region: It is recommended to choose the region closest to your customers to reduce latency and improve access speed. Availability Zone: Select according to your actual needs.If you need to purchase multiple cloud virtual machines, it is recommended that you choose different availability zones to achieve disaster recovery.For more information about selectable regions and availability zones, see Regions and Availability Zones.

Category	Required/Optional	Configuration Description
Network	Required	It represents a logically isolated network space built on the cloud platform. A private network consists of at least one subnet. The system will provide a default private network and subnet for each region. If the existing private network/subnet does not meet your requirements, you can create one in the private network console.Note:Resources within the same private network communicate with each other by default.The CVM must be created within a subnet that has the same availability zone attributes as the CVM.When manually allocating private IP addresses, the number of editable IPs in the IP input box is related to the quantities of CIDR in the subnet.
Instance	Required	Based on different underlying hardware, the cloud platform currently offers various instance types. For optimal performance, it is recommended to use the latest generation of instance types.For more details on instances, refer to Instance Specifications.
Image	Required	The cloud platform provides public images, custom images, and shared images. You can choose based on Image Types.
System Disk	Required	Used to install the operating system. The default size is 50 GB.Different regions will affect the available types of Cloud Block Storage. Please select based on the actual prompts on the page.For more information about cloud disks, please refer to Cloud Block Storage types .
Data Disk	Optional	Used to expand the storage capacity of cloud virtual machines and provide efficient and reliable storage devices. By default, no cloud block storage data disk is added.For more information about Cloud Block Storage, please refer to Cloud Block Storage types.
Scheduled Snapshots	Optional	The snapshot service provides regular backup of cloud block storage data to deal with risks such as virus infection and accidental data deletion.

Category	Required/Optional	Configuration Description
Public IP	Required	To assign a public IP address to an instance, you need to select Allocate now. The IP address assigned in this manner cannot be unbound from the instance, but it can be converted into an Elastic Public IP (EIP) for unbinding. For more information about EIP , please refer to EIP .

3.Click Next: Configuring Security Groups and Host to enter the configuration page. 4.Follow the on-page prompts to configure the following information:

Category	Required/Optional	Configuration Description
Security Group	Required	If you do not have a usable security group, you can choose Create Security Group.If you already have a usable security group, you can choose Existing security group.For more information about security groups, please refer to Security Groups.
Instance Name	Optional	User-customized, indicating the name of the cloud virtual machine to be created.If you do not define an instance name, the instance name after creation will be ”Unnamed”.If an instance name is defined, it must be within 60 characters.Note: This name is only displayed in the console and is not the hostname of the CVM.
Log in Method	Required	Set the method for users to log in to the cloud virtual machine according to actual needs.Set password: Customize the password for logging in to an instance.Associate Key Immediately (Linux Instances Only): Associate an SSH key for a more secure login to the CVM. If you do not have a key or the existing key is not suitable, you can click Create Now to create one. For more information on SSH Key, please see SSH Key.Random Password: The automatically generated password will be sent through the message center.
Security Reinforcement	Optional	Helping Users Build Server Security Systems to Prevent Data Leaks.

Category	Required/Optional	Configuration Description
Cloud Monitor	Optional	Cloud service monitoring is enabled free of charge by default. Install components to obtain host monitoring metrics and display them in the form of monitoring icons. It also supports customizing alarm thresholds. It also provides three-dimensional cloud virtual machine data monitoring, intelligent data analysis, real-time fault alarms, and personalized data report configuration, allowing users to accurately control the health of their business and cloud virtual machines.
Advanced Settings	Optional	Instances can be further configured according to actual needs.□Node name: Users can customize the computer name inside the cloud virtual machine operating system. After the cloud virtual machine is successfully put into operation, you can check it by logging in to the cloud virtual machine.□Placement group: Instances can be added to a placement group as needed to improve business availability. For detailed setup, refer to Placement Group.□Custom data: Specify custom data to configure an instance, that is, run the configured script when the instance starts. If more than one CVM is purchased at a time, custom data runs on all of them. The Linux operating system supports the Shell format, and the Windows operating system supports the PowerShell format, up to 16KB of raw data. For more details, refer to Custom Data.Note: Custom data configuration is only supported by certain public images with Cloudinit services. For specifics, refer to Cloud-Init.

5.Click Next: Confirm Configuration Information to enter the Confirm configurations page. 6.Click Enable to complete creation.

Creating an instance through an image Overview You can conveniently create cloud server instances with the same operating system, applications, and data using custom images, improving work efficiency or delivery speed. This article guides you on how to create instances using custom images. Prerequisites A custom image has been created in the account and region where you want to create the instance. If you haven't created a custom image yet, consider the following options:

Image holding status	Operational plan
Holding images locally or on other platforms	Import the system disk image file of local servers or other platforms into the cloud server's custom images. For specific operations, see Overview of Image Import.

Image holding status	Operational plan
No custom images but instances available for templates	Specific operations are described in Creating Custom Images.
Holding custom images in other regions	Copy the custom image to the region where you want to create the instance. For detailed operations, see Copying Images.
Holding custom images in other accounts	Share the custom image with the account that needs to create an instance. For specific operations, see Sharing Custom Images.

Directions 1.Log in to the Cloud Server Console. 2.In the left navigation bar, click on Image and enter the Image Management page. 3.In the status bar at the top of the Image Management page, select the region. 4.Based on the source of the image, choose a tab to enter the list interface for images. o Public Image tab: Proceed to the public image interface. o Custom Image tab: Proceed to the custom image interface. o Shared Image tab: Proceed to the shared image interface. 5.Find the image you want to use and click Create Instance in the operation column. 6.Click OK in the pop–up prompt box. 7.Follow the prompts on the screen to configure instance information and complete the creation of the instance. The region and image information is automatically filled in; please configure other instance information according to your needs. Detailed information is provided in Creating Instances through Cloud Server Application Page. Note If the custom image you select contains one or more data disk snapshots, the system will automatically create the same number of cloud hard disks as data disks based on these snapshots, with each cloud hard disk having the same capacity as its corresponding snapshot. You can increase the capacity of the cloud hard disks but cannot decrease it. Related documents You may also use the Create Instance “RunInstances” API interface to create instances using custom images. Note If you’re creating an instance from a whole machine image, first call the “DescribeImages” interface to obtain the snapshot ID associated with the image and then pass in the snapshot ID parameter when calling the Create Instance “RunInstances” interface. Otherwise, the created cloud hard disk and the corresponding snapshot ID won’t match, making it impossible to roll back snapshot data, leaving no data on the data disk and preventing normal mounting.

Batch sequential naming or specifying pattern string naming

Last Updated At: 2025-08-14 09:26:49

Overview During the creation of multiple instances, if you want the instance names/hostnames to follow certain rules, we provide automatic suffix number increment features and specified pattern string functions after batch creation. You can achieve this through both the application page and cloud API methods. • When you need to apply for n instances and hope to generate instance names/hostnames similar to ”CVM + serial number” (i.e., instance names/hostnames like CVM1, CVM2, CVM3, etc.), you can achieve this through automatic suffix number increment. • When you need to create n instances with specified instance names/hostnames that include sequential numbers starting from x, you can achieve this by specifying a single pattern string. • When you want to create n instances with multiple prefixes where each prefix specifies its own sequence number for the instance name/hostname, you can achieve this by specifying multiple pattern strings. Scope of application This document applies to setting instance names and hostnames. Operating steps Note The following operating steps are illustrated using the setting of instance names as an example. The operating steps for setting hostnames may differ slightly depending on the type of name being set. Automatic suffix number increment You can set the names of instances applied in bulk so that they have the same prefix with only the serial numbers increasing. Attention Successfully created instances default serial numbers start from 1 and cannot specify the starting serial number. The following operations are illustrated using the case where you applied for 3 instances and hope to generate instance names ”CVM + serial number” (i.e., instance names CVM1, CVM2, and CVM3). 1.Refer to the process of creating an instance application for 3 instances, and fill in the instance name according to the naming rule of ”prefix + serial number” in the ”Set Network and Host”

section, i.e., enter the instance name as CVM. As shown in the figure below:
http://imgxxfb.yun.ccb.com/raw/ce041edd4da3e85f4e0ce522d78eb077.png

2.Follow the page prompts step-by-step and complete payment. Specifying Pattern Strings Instances applied in bulk can be set with complex and specified serial number instance names. Instance names support specifying a single or multiple pattern strings; when setting the instance name, please set according to actual needs. Naming for Specifying Pattern Strings: {R:x}, where x represents the initial serial number of the generated instance name. Specifying a Single Pattern String The following operations are illustrated using the case where you need to create 3 instances and specify that the instance’s serial numbers start from 3 and increment as an example. 1.Refer to creating an instance application process and fill in the instance name according to the naming rule of ”prefix + specified pattern string {R:x}” in the

”Set Network and Host” section, i.e., enter the instance name as CVM{R:3}. As shown in the figure below: http://imgxxfb.yun.ccb.com/raw/9c7fd6f478b092ce20477b94f5ba0489.png

2.Follow the page prompts step–by–step and complete payment. Specifying Multiple Pattern Strings The following operations are illustrated using the case where you need to create 3 instances and hope that the generated instance names contain cvm, Big, and test prefixes, with serial numbers after the cvm and Big prefixes starting from 13 and 2 respectively (i.e., instance names will be cvm13–Big2–test, cvm14–Big3–test, cvm15–Big4–test) as an example. 1.Refer to creating an instance application for 3 instances and fill in the instance name according to the naming rule of ”prefix + specified pattern string {R:x}–prefix + specified pattern string {R:x}–prefix” in the ”Set Network and Host” section, i.e., enter the instance name

as cvm{R:13}–Big{R:2}–test. As shown in the figure below: http://imgxxfb.yun.ccb.com/raw/1b62106a71a2febff023fed83bcd8da6.png

2.Follow the page prompts step–by–step and complete payment. Verification Functionality After achieving bulk creation of instances through automatic suffix number increment or specifying pattern strings, you can verify it through the following operations. Verifying Setting Instance Names Log in to the cloud server console to view newly created instances; you will find that instances applied in bulk are named according to the rules you set. Verifying Setting Hostnames

1.Restart and log in to the cloud server instance. 2.In the operating system interface, execute the following command: hostname 3.Check the return results of the hostname command. If the results are similar to the following, it indicates that the setting was successful. cvm13–Big2–test 4.Repeat steps 1 through step 3 to verify other instances applied in bulk one by one.

Adjusting Instance Configuration

Last Updated At: 2025-08-14 09:26:49

Overview The hardware equipment of the instance can be adjusted quickly and easily, which is an important manifestation of the flexibility of CVM. This document describes how to upgrade and downgrade configurations and related precautions. Preconditions Adjustment of configuration operations can be conducted on instances while they are in a powered-off state. 1.log in to CVM Console, and enter the CVM List. 2.In the operation column on the right side of the instance that needs adjustment, click Action> CVM Configuration> Adjust Configuration. 3.In the ”Power Off Prompt,” carefully read the different text prompts based on the running status of the instance. • If the current instance is running, you need to carefully read the text prompt and check ”Agree to Force Power Off.” • If your instance is already powered off, you will be informed again. 4.Click Start Adjustment.

Viewing Information

Last Updated At: 2025-08-14 09:26:49

Viewing Instance Information Overview To facilitate users in viewing cloud server instance information, the cloud platform provides three paths for checking: • On the dashboard overview page, view the total number of cloud server instances under your account, along with their operating status; also see resource quantities and quotas by region. • On the cloud server page of the console, check the information of all cloud server instances in a specific region. • On the instance details page, view detailed information about a particular cloud server instance. Prerequisites You have logged in to the Cloud Virtual Machine console. Directions View instance overview information. In the left navigation bar, select Overview to enter the cloud server overview page. On this page, you can view the following information and perform the following operations: • Cloud server status, i.e., total number of cloud servers, quantity of instances expiring within 7 days, quantity of instances in recycle bin, and normal server count. • List of cloud servers awaiting renewal, and you can renew cloud servers on this page. • Quantity of resources and quotas, where you can view information about pay-as-you-go cloud servers, custom images, and snapshot capacity in each region, and apply for quotas on this page. • Cross-region search for cloud resources. Viewing Cloud Virtual Machine List Information In the left navigation bar, select Instances to enter the CVM page. On this page, you can view the following information: ID/host name, monitoring, status, AZ, architecture, host type, configuration, primary IPv4 address, billing mode, placement group, creation time, Action, etc. View Instance Monitoring Information Overview In order to facilitate users viewing monitoring information for cloud server instances, the cloud platform provides two paths for viewing as follows: • Check the monitoring information of a specific cloud server instance in the Cloud Monitor console. 1.Log in to the Cloud Server – Basic Monitoring Console. 2.Click on the instance ID to enter the monitoring information page for this instance, where you can view CVM instance monitoring information including CPU, memory, intranet bandwidth, internet bandwidth, and disk usage. • View the monitoring information of a specific cloud server instance on the instance details page in the Cloud Server console. 1.Log in to the Cloud Server console. 2.Select the region where the instance whose monitoring information is to be viewed belongs. 3.Click on the instance ID to enter the detailed page for this instance. 4.Choose the Monitoring tab to enter the monitoring information page, where you can view the monitoring information for the cloud server instance's CPU, memory, intranet bandwidth, internet bandwidth, and disk usage. View Instance Metadata Instance metadata refers to data related to the instance, which can be used to configure or manage running instances. Querying Instance Metadata Within the instance, you can access data such as local IP, public IP through instance metadata to manage connections with external applications. Taking Nanhu self-use area as an example, to view all categories of instance metadata from within a running instance, please use the following URI: http://metadata.yun.com/latest/meta-data/ You can access metadata using the cURL tool or an HTTP GET request, for example: curl http://metadata.yun.com/latest/meta-data/ • For non-existent resources, it will return HTTP error code

404 – Not Found. • All operations on instance metadata can only be performed from within the instance. Please complete the instance login operation first. For more information about logging into instances, refer to Logging into Windows Instances and Logging into Linux Instances guides. Metadata Query Example The following example illustrates how to obtain metadata version information.

1.0 2017–09–19 latest meta–data The following example shows how to view the root directory of metadata. Words ending with / represent directories, words not ending with / represent accessible data. The meanings of specific accessible data are referenced in the previous section on instance metadata classification.

instance–id instance–name local–ipv4 mac network/ placement/ public–ipv4 uuid

Image

Image

Last Updated At: 2025-08-14 09:25:05

Creating Custom Image Overview When you create an image, you can choose to launch an instance based on a public image and connect the image to your instance to deploy the software environment yourself. When the instance is running normally, you can create a new custom image based on it according to your actual needs. After you create a new custom image, you can use the image to launch more new instances with the same custom items as the original instance. Notes □Each region currently supports 50 custom images. □The following directories/files will be cleared: ○/var/log/ ○/root/.bash_history, /home/ubuntu/.bash_history (Ubuntu system) □When you create a custom image for a Linux instance, make sure that /etc/fstab does not contain data disk configuration. Otherwise, the instance created using the image will not be able to start normally. If the Linux instance for which you created the custom image has a data disk mounted, you need to comment or delete the configuration related to the self-configured data disk in /etc/fstab. □The creation process will take ten minutes or longer. The specific time depends on the data size of the instance. Make relevant preparations in advance to avoid affecting your business. Directions Creating from Instance using Console 1.In the row of the instance, click Operate > Create image. 2.In the popped-up Create Custom Image window, configure according to the following details: ○Image name and description: set your own name and description. ○Tags: add tags as needed for resource categorization, searching, and aggregation. For more information, refer to the Tag Guide. ○Create system disk image only: this option is not available if your instance has only a system disk. If there are data disks, you can select based on need. ◦ If selected, create an image of just the system disk for the instance. ◦ If unselected, and the instance includes data disks, snapshots of these will also be created simultaneously. 3.Click Confirm to create. You can click the Image in the left navigation bar to enter the Image Management Page. Using a custom image to create an instance (optional step) In the Image List, select the image you created, click Create Cloud Host to purchase a server with the same image as before. Sharing Custom Images Overview Shared image is a custom image that you have created to share with other users. Users can easily obtain shared images and get the necessary components from other users, and add their custom content. Notes □Each image can be shared with up to 50 users. □The name and description of a shared image cannot be changed and it can only be used to create the Cloud Virtual Machine instance. □Images shared with other users do not occupy their own image quota. □Images shared with other users can be deleted, but all sharing of the image must be canceled first. For details on unshare operation, see Unshare Custom Image. The obtained shared image cannot be deleted. □Images can be shared to accounts in the same region. If you need to share to different regions, you need to copy the image to different regions before sharing. □The obtained shared image cannot be shared with other users. Directions Obtain Account ID The shared image is identified by

the unique ID of the peer account. You can notify the user to obtain it in the following ways: 1.Log in to the CVM Console. 2.Click the account name in the upper right corner and select Account Center. 3.On the Account Information management page, view and record the account ID. 4.Inform the other party to send the obtained account ID to you. Share via the Console 1.Log in to the CVM Console. In the left–hand navigation bar, click Image. 2.Select the Custom Image tab, enter custom image management page. 3.In the custom image list, select the custom image you want to share, and click Share on the right. 4.In the pop–up Shared Image window, enter the other party's account ID, and click Share. 5.Notify the other party to log in to CVM console, and select Image > Shared Images to view the shared images. To share with multiple users, repeat the above steps. Unsharing Custom Image Overview This document guides users to unshare custom images. Users can terminate the sharing status of the image shared with others at any time to decide not to share it with other users. This operation does not affect instances that other users have already created using this shared image, but other users will no longer be able to view this image or create more instances using this image. Directions Unsharing via the Console 1.Log in to the CVM Console. 2.In the left–hand navigation bar, click Image. 3.Select the Self–Definition Image tab, in the Self–Definition Image list, select the image you want to stop sharing, then click More > Unshare. 4.In the new page, select the unique ID of the account you want to unshare from, click Unshare. 5.In the pop–up dialog box, click Confirm to complete the unsharing of the image. Deleting Custom Image Overview This document guides you to delete custom images. Notes Before deleting, note the following: ●After you delete custom images, you cannot create instances using this image. However, this does not affect instances that have already been started. ●A shared image cannot be deleted. You must first unshare all images before deleting them. To unshare images, see Unshare Custom Image. ●Only custom images can be deleted. Public images and shared images cannot be deleted actively, and obtained images cannot be deleted either. Directions 1.Log in to the CVM Console. 2.In the left navigation bar, click Image, select the Custom Image tab to enter the custom image management page. 3.Select the operation method for deleting a custom image based on actual needs. –Delete a single image: Locate the custom Definition image to be deleted in the list, click More > Delete. –Deleting multiple images: Select all custom images to be deleted in the list, click Delete at the top. 4.In the popped–up prompt box, click OK. If the deletion fails, possible reasons will be prompted. Importing Image Overview In addition to using the Create Custom Image feature, Cloud Platform also supports to import. You can import image files of local server system disk or other platform into the custom image of Cloud Virtual Machine (CVM). After you import, you can use the imported image to create the Cloud Virtual Machine or reinstall the system on an existing Cloud Virtual Machine. Import Preparation You need to prepare an image file that meets the import restrictions in advance. Image limits of Linux system type:

Image Attribute	Condition
Operating System	●Images based on CentOS, Ubuntu, Debian, CoreOS, openSUSE, and SUSE distributions.●Supports 32–bit and 64–bit

Image Attribute	Condition
Image Format	●Support RAW, VHD, QCOW2, VMDK image formats●Use qemu–img info imageName
File System Type	GPT partition unsupported
Image Size	Use qemu–img info imageName
Network	●By default, eth0 network API is provided for the instance.●IPv6 is not supported yet.●Users can query the network configuration of an instance through the metadata service within the instance. For details, see Instance Metadata.
Drive	●The image must be installed with the Virtio driver of the virtualization platform KVM. For details, see Import Image on Linux to Check Virtio Driver.●It is recommended to install cloudinit in the image. For details, see Import Image on Linux to Install Cloudinit.●If the image cannot be installed with cloudinit due to other reasons, use Forced Import Image to self–configure instance.
Kernel Limits	The image should preferably be a native kernel. Modification may cause the Cloud Virtual Machine to fail to import.

Import Steps 1.Log in to the CVM Console. Click the Image in the left navigation bar. 2.Select Custom Image, Click Import Image. 3.According to the requirements of the operation interface, first enable the Cloud Object Storage service, then create a bucket, upload the image file to the bucket and obtain the image file URL. 4.Click Next. 5.Fill out the form according to the actual situation, Click Start Importing. Import Failed After the image import operation is performed on the console, the task may fail due to some reasons. If the task fails, you can troubleshoot based on the following information. Troubleshooting Cause of Failure For detailed error prompts and error descriptions, see the error code. InvalidUrl: The Cloud Object Storage Link is Invalid The InvalidUrl error is reported, error prompts: An incorrect Cloud Object Storage link was entered on the image import page. Possible reasons are as follows: The image link entered is not a Cloud Object Storage service. The object address of Cloud Object Storage does not have public read and private write permissions. The access permission to the Cloud Object Storage file is private read, but the signature has expired. A Cloud Object Storage link in another region was entered. The user's image file has been deleted. If you receive an error reporting that the Cloud Object Storage link is invalid, you can troubleshoot the problem based on the reasons listed above. InvalidFormatSize: The format or Size does not Meet the Requirements The InvalidFormatSize error is reported, error prompts: The format or size of the pre–imported image does not meet the limits of the import image feature. The

limits are as follows: □Import image supports 4 image file formats: qcow2, vhd, vmdk, and raw. □The actual file size of the imported image shall not exceed 50GB (based on the image file converted to qcow2 format). □The size of the system disk of the import image must not exceed 500 GB. After an error reporting that the format or size does not meet the requirements is received: □Based on the image format conversion content of Linux Image Creation, you can convert the image file to a suitable file format, condense the image content to meet the size limit, and then re-import the image. VirtioNotInstall: Virtio Driver is not Installed The VirtioNotInstall error is reported, error prompts: The Virtio driver is not installed on the pre-imported image. Cloud Platform uses KVM virtualization technology, which requires the virtio driver to be installed in the image that the user imports. The Virtio driver has been installed on most Linux operating systems except a few user-customized Linux operating systems. For Windows operating system, you need to manually install the Virtio driver: □For Linux image import, see the document Check Virtio Driver in Linux System. □For Windows image import, see the document Windows Image Creation to install the Virtio driver. CloudInitNotInstalled: The Cloud-init Program is not Installed The CloudInitNotInstalled error is reported, error prompts: The cloud-init driver is not installed on the pre-imported image. Cloud Platform uses the open source program cloud-init to initialize the sub-machine. Therefore, if the cloud-init program is not installed, the user's sub-machine initialization will fail. □For Linux image import, see the document Install cloud-init on Linux System. □After cloud-init/cloudbase-init is installed, replace the configuration file according to the documentation so that the sub-machine can pull data from the correct data source when it starts. PartitionNotPresent: Partition Information is Lost The PartitionNotPresent error is reported, error prompts: The imported image is incomplete. Check whether the boot partition is included when making the image. RootPartitionNotFound: The Root Partition is Lost The RootPartitionNotFound error is reported, error prompts: The imported image is not detected to contain a root partition. Check the image file. The reasons that have occurred are as follows for reference: □The installation package file is uploaded. □The data disk image is uploaded. □The boot partition image is uploaded. □The wrong file is uploaded. InternalError: Unknown Error The InternalError error is reported, error prompts: The import image service does not include the error cause. Contact customer service to handle this problem, and the technician will solve the problem as soon as possible. Error Code

Error Code	Error Cause	Suggested Handling Method
InvalidUrl	COS link is invalid.	Check if the COS link is the same as the import image link.
InvalidFormatSize	The format or size does not meet the requirements.	The image must meet the requirements of Image Format and Image Size in the import preparation.
VirtioNotInstall	Virtio driver is not installed.	The virtio driver needs to be installed for images. For details, see Driver section in the import preparation.

Error Code	Error Cause	Suggested Handling Method
PartitionNotPresent	Partition information is not found.	The image is damaged, which may be caused by the wrong image creation method.
CloudInitNotInstalled	cloud-init is not installed.	The cloud-init needs to be installed for Linux images. For details, see Driver section in the import preparation.
RootPartitionNotFound	Root partition is not detected.	The image is damaged, which may be caused by the wrong image creation method.
InternalError	Other Errors	Contact customer service to handle the problem.

Checking Virtio Driver in Linux System The kernel of Cloud Virtual Machine system needs to support virtio drivers (including block device driver virtio_blk and network interface driver virtio_net) to run properly on Cloud Platform. For virtio_blk driver which is not compiled into the kernel, it also needs to be included in initramfs (or initrd) file, then the Cloud Virtual Machine can work properly. This document takes the CentOS operating system as an example to guide you on how to check and repair the support for the Virtio driver in the image before importing it. Directions Step 1: Check whether the kernel supports Virtio driver Execute the following command to check whether the current kernel supports the Virtio driver. grep –i virtio /boot/config–$(uname –r) The result resembles the following is returned:

```
[root@VM_0_120_centos ~]# grep -i virtio /boot/config-$(uname -r)
CONFIG_VIRTIO_VSOCKETS=m
CONFIG_VIRTIO_VSOCKETS_COMMON=m
CONFIG_VIRTIO_BLK=m
CONFIG_SCSI_VIRTIO=m
CONFIG_VIRTIO_NET=m
CONFIG_VIRTIO_CONSOLE=m
CONFIG_HW_RANDOM_VIRTIO=m
CONFIG_DRM_VIRTIO_GPU=m
CONFIG_VIRTIO=m
# Virtio drivers
CONFIG_VIRTIO_PCI=m
CONFIG_VIRTIO_PCI_LEGACY=y
CONFIG_VIRTIO_BALLOON=m
CONFIG_VIRTIO_INPUT=m
# CONFIG_VIRTIO_MMIO is not set
```

If the kernel supports the virtio driver (both virtio_blk and virtio_net is supported), and virtio_blk driver is compiled into the kernel (i.e. CONFIG_VIRTIO_BLK=y), the kernel supports import and no subsequent confirmation is required. If the virtio_blk driver is compiled into a kernel module (i.e. CONFIG_VIRTIO_BLK=m), you need to continue the subsequent confirmation steps to confirm the virtio_blk driver is correctly included in the initramfs (or initrd) file. If in the returned result, the values of

CONFIG_VIRTIO_BLK parameter and the CONFIG_VIRTIO_NET parameter are m, proceed to Step 2. If in the returned result, the values of CONFIG_VIRTIO_BLK parameter and CONFIG_VIRTIO_NET parameter are y, it indicates that the operating system includes the Virtio driver and you can directly import the custom image. If in the returned result, there is no information about CONFIG_VIRTIO_BLK parameters and CONFIG_VIRTIO_NET parameter, it indicates that the operating system does not support import. Step 2: Check whether the temporary file system contains the Virtio driver If in the execution result of Step 1, parameter value is m, you need to further check and confirm whether the temporary file system initramfs or initrd includes virtio drive. Execute the corresponding command according to the operating systems: CentOS 6/CentOS 7/CentOS 8/RedHat 6/RedHat 7 operating systems: lsinitrd /boot/initramfs-$(uname -r).img | grep virtio RedHat 5/CentOS 5 operating systems: mkdir -p /tmp/initrd && cd /tmp/initrd zcat /boot/initrd-$(uname -r).img | cpio -idmv find . -name "virtio*" Debian/Ubuntu operating systems: lsinitramfs /boot/initrd.img-$(uname -r) | grep virtio The result resembles the following is returned:

```
[root@VM_0_120_centos ~]# lsinitrd /boot/initramfs-$(uname -r).img | grep virtio
-rw-r--r--   1 root     root         7744 Apr 21  2018 usr/lib/modules/3.10.0-862.el7.x86_64/kernel/drivers/block/virtio_blk.ko.xz
-rw-r--r--   1 root     root        12944 Apr 21  2018 usr/lib/modules/3.10.0-862.el7.x86_64/kernel/drivers/char/virtio_console.ko.xz
-rw-r--r--   1 root     root        14296 Apr 21  2018 usr/lib/modules/3.10.0-862.el7.x86_64/kernel/drivers/net/virtio_net.ko.xz
-rw-r--r--   1 root     root         8176 Apr 21  2018 usr/lib/modules/3.10.0-862.el7.x86_64/kernel/drivers/scsi/virtio_scsi.ko.xz
drwxr-xr-x   2 root     root            0 Jan 21  2019 usr/lib/modules/3.10.0-862.el7.x86_64/kernel/drivers/virtio
-rw-r--r--   1 root     root         4556 Apr 21  2018 usr/lib/modules/3.10.0-862.el7.x86_64/kernel/drivers/virtio/virtio.ko.xz
-rw-r--r--   1 root     root         9664 Apr 21  2018 usr/lib/modules/3.10.0-862.el7.x86_64/kernel/drivers/virtio/virtio_pci.ko.xz
-rw-r--r--   1 root     root         8280 Apr 21  2018 usr/lib/modules/3.10.0-862.el7.x86_64/kernel/drivers/virtio/virtio_ring.ko.xz
```

It can be seen that initramfs has already included virtio_blk driver, and its dependencies virtio.ko, virtio_pci.ko and virtio_ring.ko, you can directly import a custom image. If the initramfs or initrd does not include the virtio driver, proceed to Step 3. Step 3: Reconfigure the temporary file system If the execution result of Step 2 shows temporary file system initramfs or initrd is not included virtio driver, you need to reconfigure the temporary file system initramfs or initrd to contain virtio drive. Select the corresponding operation according to the operating system: CentOS 8/RedHat 8 操作系统： mkinitrd -f --allow-missing --with=virtio_blk --preload=virtio_blk --with=virtio_net --preload=virtio_net --with=virtio_console --preload=virtio_console /boot/initramfs-$(uname -r).img $(uname -r) CentOS 6/CentOS 7/RedHat 6/RedHat 7 operating systems: mkinitrd -f --allow-missing --with=xen-blkfront --preload=xen-blkfront --with=virtio_blk --preload=virtio_blk --with=virtio_pci --preload=virtio_pci --with=virtio_console --preload=virtio_console /boot/initramfs-$(uname -r).img $(uname -r) RedHat 5/CentOS 5 operating systems: mkinitrd -f --allow-missing --with=xen-vbd --preload=xen-vbd --with=xen-platform-pci --preload=xen-platform-pci --with=virtio_blk --preload=virtio_blk --with=virtio_pci --preload=virtio_pci --with=virtio_console --preload=virtio_console /boot/initrd-$(uname -r).img $(uname -r) Debian/Ubuntu operating systems: echo -e 'xen-blkfront\nvirtio_blk\nvirtio_pci\nvirtio_console' >> /etc/initramfs-tools/modules mkinitramfs -o /boot/initrd.img-$(uname -r) Appendix Downloading and Compiling the Kernel Download the kernel package 1.Execute the following command to install necessary components required for compiling the kernel. yum install -y ncurses-devel gcc make wget 2.Execute the following command to query the current kernel version used by the system. uname -r The query returns results similar to the following,

indicating that the current kernel version being used is 2.6.32–642.6.2.el6.x86_64.

```
[root@VM_0_139_centos ~]# uname -r
2.6.32-642.6.2.el6.x86_64
```

3.Go to the Linux Kernel download page and download the corresponding kernel source code for the version used. For instance, for the kernel version 2.6.32–642.6.2.el6.x86_64, download the package linux–2.6.32.tar.gz, which is available at https://mirrors.edge.kernel.org/pub/linux/kernel/v2.6/linux–2.6.32.tar.gz. 4.Execute the following command to switch directories. cd /usr/src/ 5.Execute the following command to download the package. wget https://mirrors.edge.kernel.org/pub/linux/kernel/v2.6/linux–2.6.32.tar.gz 6.Execute the following command to extract the package. tar –xzf linux–2.6.32.tar.gz 7.Execute the following command to create a symbolic link. ln –s linux–2.6.32 linux 8.Execute the following command to switch directories. cd /usr/src/linux Compile the kernel 1.Execute the following commands in sequence to compile the kernel. make mrproper cp /boot/config–$(uname –r)./.config make menuconfig Enter the ”Linux Kernel vX.X.XX Configuration” interface. See the figure below.

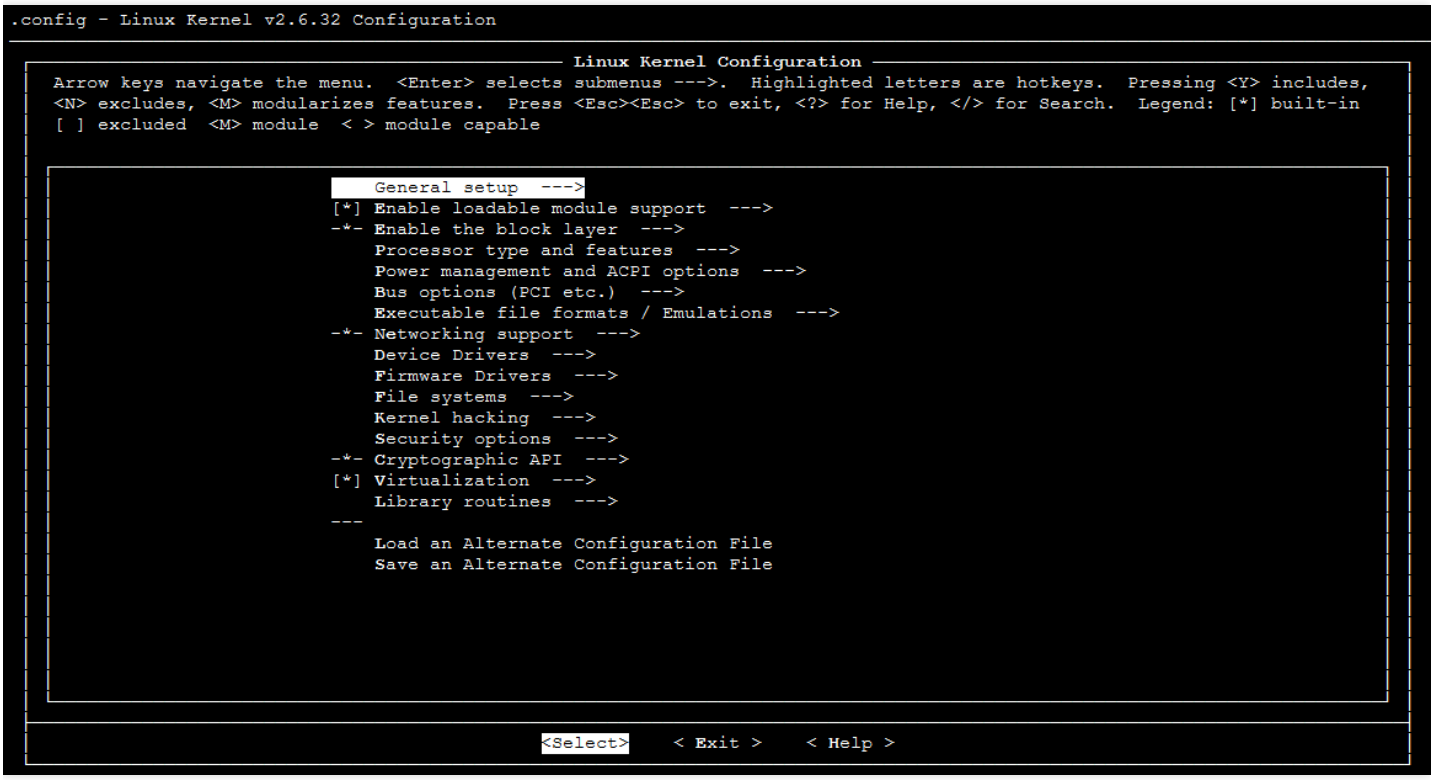

Note If you do not enter the ”Linux Kernel vX.X.XX Configuration” interface, execute Step 18. The ”Linux Kernel vX.X.XX Configuration” interface: o Press the ”Tab” key or the direction keys ”↑” ”↓” to move the cursor. o Press ”Enter” to select or execute the item highlighted by the cursor. o Press the space bar to select the item highlighted by the cursor; ”*” indicates compilation into the kernel, and ”M” indicates compilation as a module. 2. Press the ”↓” key to move the cursor to ”Virtualization” and press the space

bar to select ”Virtualization”. 3. Press ”Enter” at ”Virtualization” to enter the detailed Virtualization interface. 4. In the detailed Virtualization interface, check whether the Kernel-based Virtual Machine (KVM) support option is selected. See the figure below.

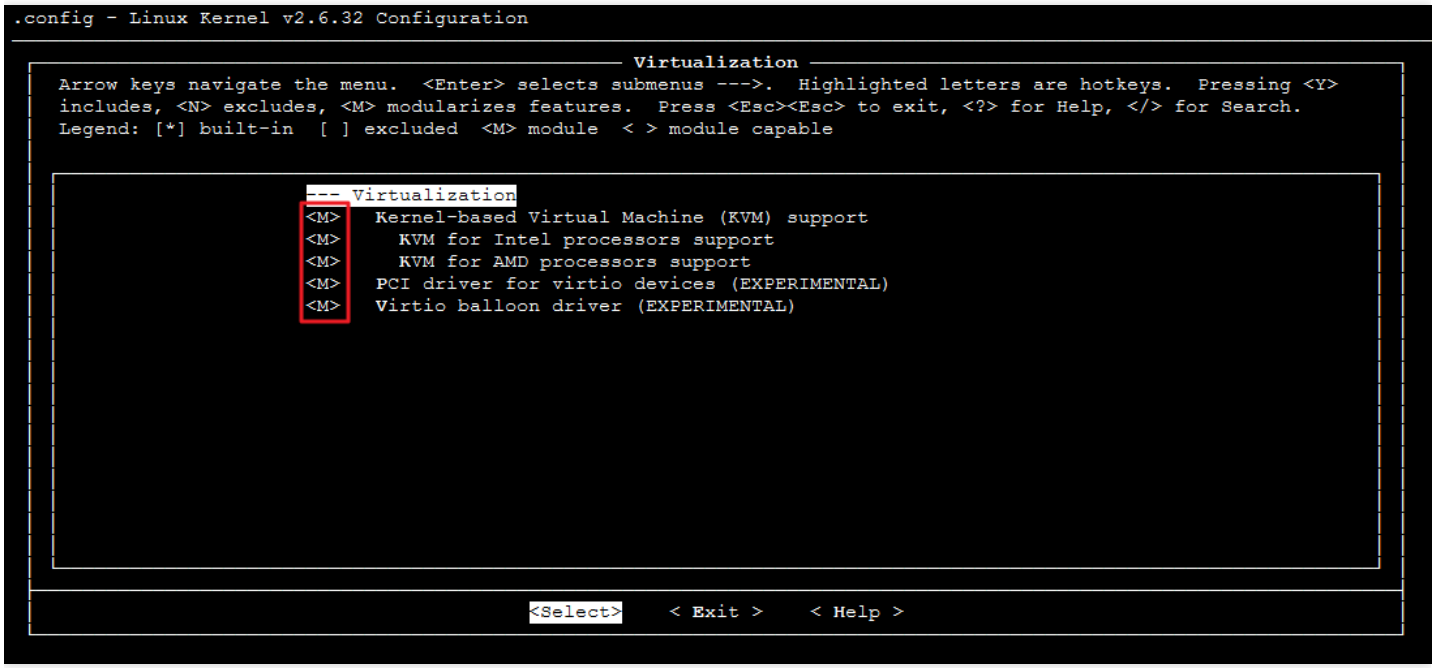

If not selected, press the space bar to select the ”Kernel-based Virtual Machine (KVM) support” option. 5. Press ”Esc” to return to the main ”Linux Kernel vX.X.XX Configuration” interface. 6. Press the ”↓” key to move the cursor to ”Processor type and features” and press ”Enter” to enter the detailed Processor type and features interface. 7. Press the ”↓” key to move the cursor to ”Paravirtualized guest support” and press ”Enter” to enter the detailed Paravirtualized guest support interface. 8. In the detailed Paravirtualized guest support interface, check whether ”KVM paravirtualized clock” and ”KVM Guest support” are selected. See the figure below:

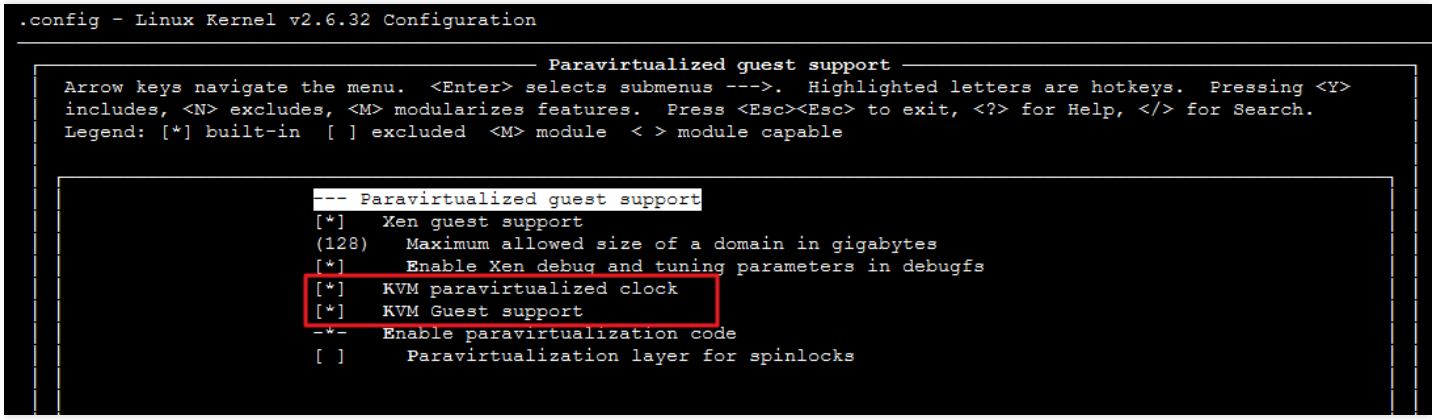

If not selected, press the space bar to select the ”KVM paravirtualized clock” and ”KVM Guest support” options. 9. Press ”Esc” to return to the main ”Linux Kernel vX.X.XX Configuration” interface. 10. Press the

”↓” key to move the cursor to ”Device Drivers” and press ”Enter” to enter the detailed Device Drivers interface. 11. Press the ”↓” key to move the cursor to ”Block devices” and press ”Enter” to enter the detailed Block devices interface. 12. In the detailed Block devices interface, check whether ”Virtio block driver (EXPERIMENTAL)” is selected. See the figure below:

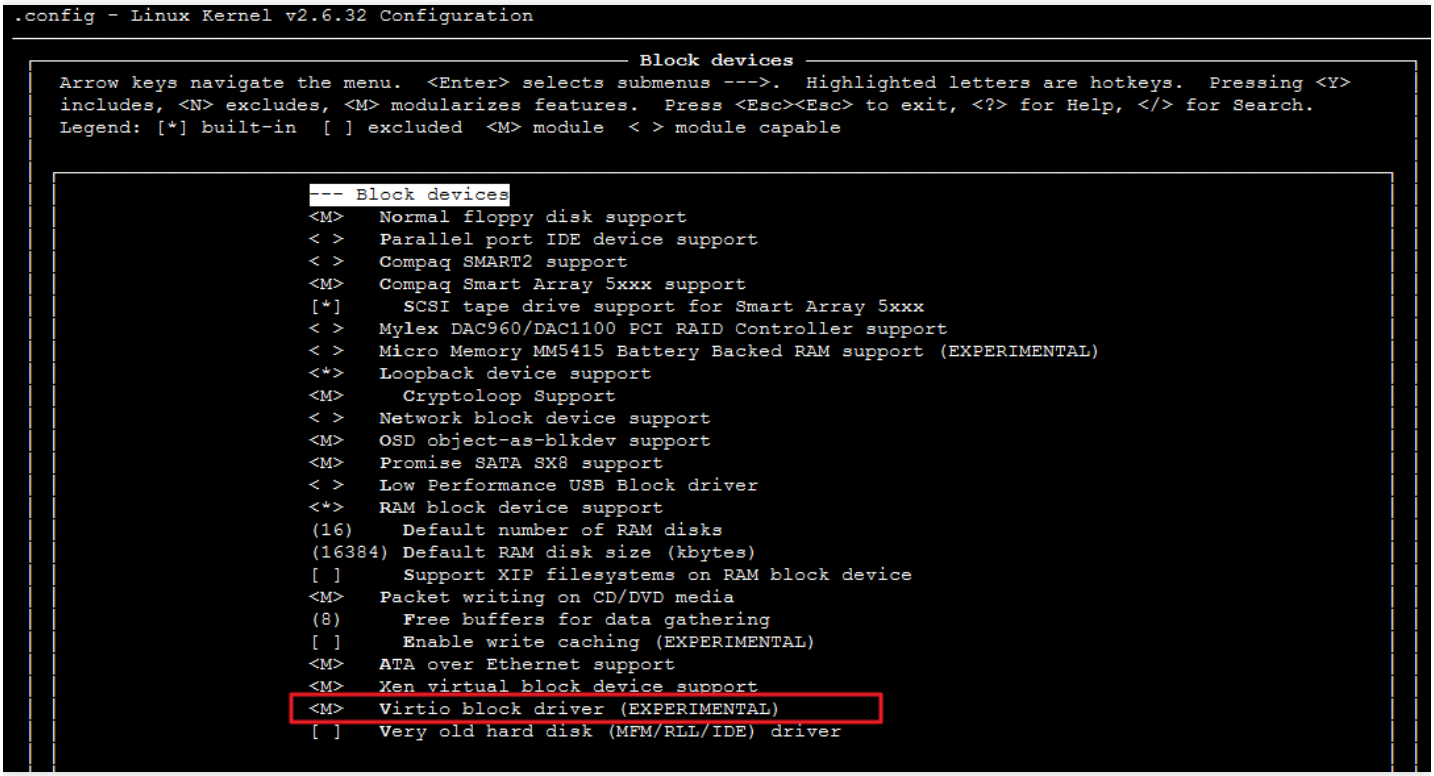

If not checked, press the spacebar to select the ”Virtio block driver (EXPERIMENTAL)” option. 13. Press ”Esc” to return to the Device Drivers detail interface. 14. Press the ”↓” key to move the cursor to ”Network device support,” then press ”Enter” to enter the Network device support detail interface. 15. In the Network device support detail interface, check whether the ”Virtio network driver (EXPERIMENTAL)” is selected. As shown in the figure below:

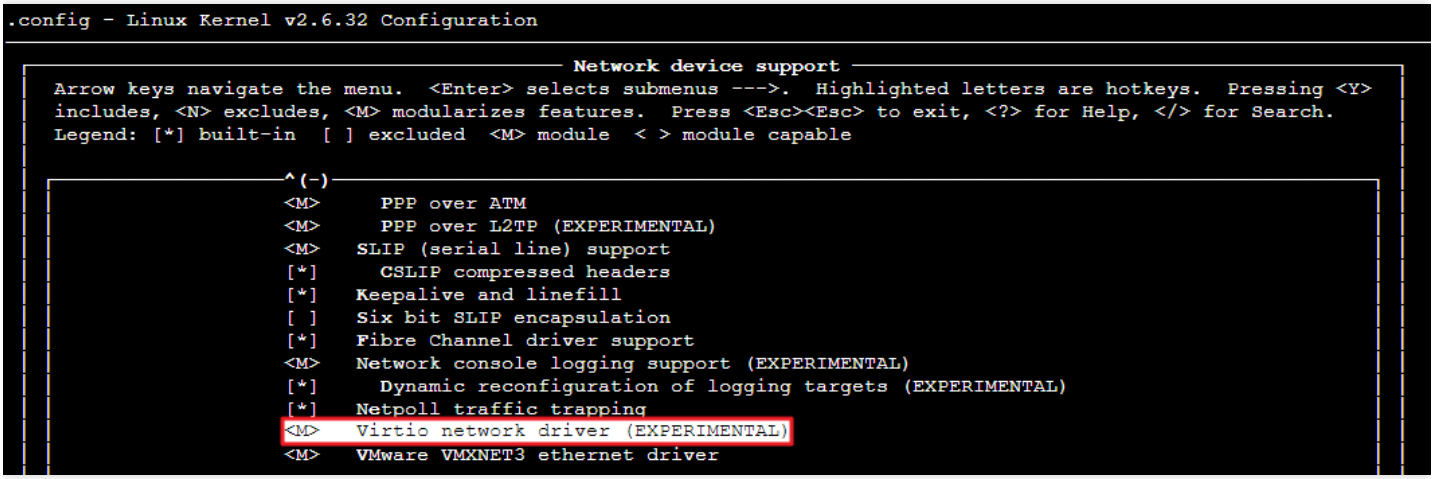

If not checked, press the spacebar to select the ”Virtio network driver (EXPERIMENTAL)” option. 16. Press ”Esc” to exit the kernel configuration interface and, according to the pop–up prompt, select ”YES” to save the .config file. 17. Refer to Step 1: Check whether the kernel supports Virtio drivers to verify whether the Virtio driver has been correctly configured. 18. (Optional) Execute the following command to manually edit the .config file. Note If any of the following conditions apply to you, it is recommended to perform this operation: o If after checking, you find that there is still no configuration information related to the Virtio driver in the kernel. o When compiling the kernel, you cannot enter the kernel configuration interface or did not successfully save the .config file. make oldconfig make prepare make scripts make make install 19. Execute the following commands in sequence to view the installation status of the Virtio driver. find /lib/modules/”(uname –r)”/ –name ”virtio.*” | *grep –E ”virtio.*”* grep –E ”virtio.*” < /lib/modules/”(uname—r)”/—name”virtio.∗”|grep—E”virtio.∗”grep—E”virtio.∗”</lib/modules/”(uname –r)”/modules.builtin If any command’s return results display file lists such as virtio_blk, virtio_pci.virtio_console, etc., it indicates that you have correctly installed the Virtio driver. Install cloud–init on Linux System Overview Cloud–init mainly can customize the configuration when an instance is initialized for the first time. If the cloud–init service is not installed for the imported image, the instance started based on the image will not be initialized normally, causing the image import to fail. This document guides you to install the cloud–init service. Notes Before a Linux system image is imported, make sure that the cloud–init service has been correctly installed within your image. Prerequisites The server with cloud–init installed can access the public network normally. Directions Downloading the Cloud–init Source Package Note: □In the case of normal installation, cloud–init–17.1 has the best compatibility with Cloud Platform, which ensures that all configuration items of the Cloud Virtual Machine created using this image can be initialized normally. It is recommended to select cloud–init–17.1.tar.gz installation version. This document takes cloud–init–17.1 as an example. wget https://launchpad.net/cloud–init/trunk/17.1/+download/cloud–init–17.1.tar.gz Install cloud–init 1.Execute the following command to decompress the cloud–init installation package. Note: If you are using Ubuntu, switch to the root account. tar –zxvf cloud–init–17.1.tar.gz 2.Execute the following command to enter the directory of the decompressed cloud–init installation package (that is, enter the cloud–init–17.1 directory). cd cloud–init–17.1 3.Based on the operating system version, install Python–pip. □For CentOS 6/7 series, execute the following command: yum install python3–pip –y □For Ubuntu series, execute the following command: apt–get –y install python3–pip 4.Execute the following command to upgrade pip. python3 –m pip install ––upgrade pip 5.Execute the following command to install the dependency package. Note: For later version of Cloud–init dependency component requests version 2.20.0, Python 2.6 has been deprecated. If the Python interpreter of the image environment is Python 2.6 or earlier, before the cloud–init dependency package is installed, you can execute pip install ’requests<2.20.0’ command to install earlier version of requests versions 2.20.0. pip install –r requirements.txt 6.Install the cloud–utils component according to the operating system version. –For CentOS 6 series, execute the following command: yum install cloud–utils–growpart dracut–modules–growroot –y dracut –f –For CentOS 7 series, execute the following command: yum install cloud–utils–growpart –y –For Ubuntu series, execute the following command: apt–

get install cloud–guest–utils –y 7.Execute the following command to install cloud–init. python setup.py build python setup.py install ––init–system systemd Note: Optional parameters for ––init–system are: (systemd, sysvinit, sysvinit_deb, sysvinit_freebsd, sysvinit_openrc, sysvinit_suse, upstart) [default: None]. Select according to the auto–start service management method used by the current operating system. If you make an incorrect selection, the cloud–init service will fail to start at boot. This document takes systemd self–start service management as an example. Modify cloud–init Configuration File 1.Download cloud.cfg according to different operating systems. 2.Replace the content of /etc/cloud/cloud.cfg with the content of the downloaded cloud.cfg file. Add syslog User Execute the following command to add a syslog user. useradd syslog Setting Cloud–init Service to Start at Boot □If the operating system uses the systemd self–start management service, execute the following command to set it. You can execute the command strings /sbin/init | grep ”/lib/system”; if there’s any returned information, then the operating system uses systemd for automatic startup service management. □For Ubuntu or Debian operating systems, execute the following command. ln –s /usr/local/bin/cloud–init /usr/bin/cloud–init □All operating systems need to execute the following command. systemctl enable cloud–init–local.service systemctl start cloud–init–local.service systemctl enable cloud–init.service systemctl start cloud–init.service systemctl enable cloud–config.service systemctl start cloud–config.service systemctl enable cloud–final.service systemctl start cloud–final.service systemctl status cloud–init–local.service systemctl status cloud–init.service systemctl status cloud–config.service systemctl status cloud–final.service □For CentOS and RedHat operating systems, execute the following command. Replace the /lib/systemd/system/cloud–init–local.service file with the following content: [Unit] Description=Initial cloud–init job (pre–networking) Wants=network–pre.target After=systemd–remount–fs.service Before=NetworkManager.service Before=network–pre.target Before=shutdown.target Conflicts=shutdown.target RequiresMountsFor=/var/lib/cloud [Service] Type=oneshot ExecStart=/usr/bin/cloud–init init ––local ExecStart=/bin/touch /run/cloud–init/network–config–ready RemainAfterExit=yes TimeoutSec=0 Output needs to appear in instance console output StandardOutput=journal+console [Install] WantedBy=cloud–init.target Replace the /lib/systemd/system/cloud–init.service file with the following content: [Unit] Description=Initial cloud–init job (metadata service crawler) Wants=cloud–init–local.service Wants=sshd–keygen.service Wants=sshd.service After=cloud–init–local.service After=systemd–networkd–wait–online.service After=networking.service After=systemd–hostnamed.service Before=network–online.target Before=sshd–keygen.service Before=sshd.service Before=systemd–user–sessions.service Conflicts=shutdown.target [Service] Type=oneshot ExecStart=/usr/bin/cloud–init init RemainAfterExit=yes TimeoutSec=0 Output needs to appear in instance console output StandardOutput=journal+console [Install] WantedBy=cloud–init.target

□If the operating system uses the sysvinit self-start management service, execute the following com mand to set it.
You can execute the command `strings /sbin/init | grep "sysvinit"`; if there's any returned information, then the operating system uses sysvinit for automatic startup service management.

chkconfig --add cloud-init-local
chkconfig --add cloud-init
chkconfig --add cloud-config
chkconfig --add cloud-final
chkconfig cloud-init-local on
chkconfig cloud-init on
chkconfig cloud-config on
chkconfig cloud-final on
Related Operations
After the following operations are complete, do not restart the server. Otherwise, you will need to re-perform the following operations.
1.Execute the following command to check whether the cloud-init configuration is successful.
cloud-init init –local
If similar information is returned, it indicates that cloud-init has been successfully configured.
Cloud-init v. 17.1 running 'init-local' at Fri, 01 Apr 2022 01:26:11 +0000. Up 38.70 seconds.
2.Execute the following command to delete the cache records for cloudinit.
rm -rf /var/lib/cloud
3.For Ubuntu or Debian operating systems, execute the following command.

4.For Ubuntu or Debian operating systems, you need to change

This file describes the network interfaces available on your system
and how to activate them. For more information, see interfaces(5).

Appendix
Solving the Problem of Unable to Install Python-pip
If you encounter errors such as the package not existing or unable to install when installing Python-pip, you can refer to the following steps based on your actual operating system in use:
• CentOS 6/7 series
1.Execute the following command to set up the EPEL repository.
yum install epel-release -y
2.Execute the following command to install Python-pip.
yum install python3-pip -y
• Ubuntu series
1.Execute the following command to clear the cache.
apt-get clean all
2.Execute the following command to update the package list.
apt-get update -y
3.Execute the following command to install Python-pip.
apt-get -y install python3-pip
Image Format Conversion
Overview
Currently Cloud Platform virtual servers support importing image files in the following formats: RAW, VHD, QCOW2, VMDK. Image files in other formats need to be converted before they can be imported. This article introduces how to use the qemu-img tool to convert image files of other formats into VHD or RAW format.
Directions
Note

This article uses Ubuntu 20.04 and CentOS 7.8 operating systems as examples for image format conversion. There are certain differences between different versions of operating systems, please refer to the documentation based on your actual situation.
Install qemu-img
1.Execute the following command to install qemu-img.
o Ubuntu:
apt-get update # Update package list
apt-get install qemu-utils # Install qemu-img tool
o CentOS:
yum install qemu-img
2.Execute the following command to convert the image format.
qemu-img convert -f qcow2 -O raw test.qcow2 test.raw
Parameter description is as follows:
o -f: The value of this parameter is the format of the source image file.
o -O (must be uppercase): Parameter values are the target image format, name of the source image file, and name of the target file. After the conversion is complete, the target file will appear in the directory where the source image file is located.
Forcibly Importing Image
Overview
When the user's Linux image cannot install cloudinit for some reason, you can use Forced Import Image feature to import the image. If the user uses the forced import image to import, Cloud Platform will not be able to initialize the configuration of the user's Cloud Virtual Machine. You need to set a script to configure the Cloud Virtual Machine according to the configuration file provided by Cloud Platform. This document guides users on how to configure the Cloud Virtual Machine under the premise of forced import image.
Cloud Platform provides a cdrom device containing configuration information for users to configure themselves. Users need to mount cdrom and read mount_point/qcloud_action/os.conf information to configure. If the user needs to use other configuration data or UserData, they can directly read the files under mount_point/.
os.conf Configuration File Content
The basic contents of os.conf are as follows:
hostname=VM_10_20_xxxx
password=GRSgae1fw9frsG.rfrF
eth0_ip_addr=10.104.62.201
eth0_mac_addr=52:54:00:E1:96:EB
eth0_netmask=255.255.192.0
eth0_gateway=10.104.0.1
dns_nameserver="10.138.224.65 10.182.20.26 10.182.24.12"
Note: The parameter names in the above information are for reference only, and the parameter values are only examples.
The meaning of each parameter in os.conf is as follows:

Parameter Name	Parameter Meaning
hostname	Host name
password	Encrypted password
eth0_ip_addr	LAN IP of the eth0 network interface
eth0_mac_addr	MAC address of the eth0 network interface

eth0_netmask	Subnet mask of the eth0 network interface
eth0_gateway	Gateway of the eth0 network interface
dns_nameserver	DNS resolution server

Limitation Factor
□The image still needs to meet the imported image limits in Import Image for Linux images (except cloudinit).
□The system partition where the image is imported is not full.
□The imported image must not have vulnerabilities that can be exploited remotely.
□It is recommended that users change the password immediately after the instance is created successfully by forced import image.
Configure Script Resolution
Notes
□The script is automatically executed at boot. Implement this requirement according to the type of operating system.
□The script must mount /dev/cdrom, and read os_action/os.conf file under the mount point to obtain configuration information.
□The password placed in the cdrom by Cloud Platform is an encrypted password. Users can use chpasswd -e to set it.
The encrypted password may contain special characters. It is recommended to put it in a file first and then use chpasswd -e < passwd_file to set it.
□When recreating an image using an instance created by forced import image, you need to ensure that the script is still executed to ensure that the instance is correctly configured. You can also install cloudinit in this instance.
Directions
1.Create the os_config script according to the following script sample.
Users can modify the os_config script based on their actual situation.

2.Place os_config script in /etc/init.d/ directory and execute the following command.
chmod +x /etc/init.d/os_config
chkconfig --add os_config
3.Execute the following command to check whether os_config has been added to the startup services.
chkconfig --list
Export Image
The Cloud Platform supports exporting custom images to the object storage COS bucket. You can export the required images through this feature.
Prerequisites:
• Currently, using this function requires an application. Please contact online customer service to apply for the activation of the function.
• Object storage services have been activated via the object storage console.
• A bucket has been created in the region where the custom image is located; details can be found under Create Bucket.
Precautions:
• Exporting Windows custom images is currently unsupported.
• The single capacity of system disks and data disks of custom images cannot exceed 500GB.

• When exporting whole machine images, the number of data disks cannot exceed five.
Directions
1. Log in to the cloud server console and select Images from the left navigation bar.
2. On the "Images" page above, select the region where the custom image to be exported is located and click on the Custom Image tab.
3. Select More > Export Image on the right side of the row where the image is located.
4. In the pop-up "Export Image" window, perform the following settings.
o COSBucket: Choose the bucket where the exported image will reside, which needs to be in the same region as the image.
o Prefix name of the exported file: Customize the prefix name of the exported file.
Check "Agree to authorize CVM to access my COSBucket."
5. Click OK to start the image export task.
6. Click OK in the confirmation window that pops up. The export time depends on the size of the image and the busyness of the task queue, so please wait patiently. After the export task is completed, the image file will be stored in the target bucket. You may go to the bucket list page, click on the bucket ID to enter the detail page, and the file named <custom prefix name>_xvda.raw is the exported image file.
Create Linux Image
Overview
This document guides you to create a Linux image.
Directions
Preparations
When you export a system disk image, you need to check the following:
Note: If you export via a data disk image, you can skip this operation.
Check OS Partition and Starting Mode
1.Execute the following command to check whether the OS partition is a GPT partition.
sudo parted -l /dev/sda | grep 'Partition Table'
-If the returned result is msdos, it is an MBR partition, then proceed to the next step.
-If the returned result is gpt, it is a GPT partition. Currently, GPT partitions are not supported for service migration.
2.Execute the following command to check whether the operating system is started in EFI mode.
sudo ls /sys/firmware/efi
-If the file exists, the current operating system is started in EFI mode.
-If the file does not exist, proceed to the next step.
Check System Key Files
The key system files that need to be checked include but are not limited to the following:
Note: Follow the standards of the relevant distribution to ensure that the location and permissions of key system files are correct and can be read and written normally.
□/etc/grub2.cfg: It is recommended to use the UUID for mounting root in kernel parameters, as other methods (such as root=/dev/sda) may result in system failure to boot. The steps for mounting are as follows:
i. Execute the following command to obtain the file system name of /root.
df –TH
The return result is shown in the figure below, indicating that the file system name of /root is /dev/vda1.

ii. Execute the following command to get the UUID.

blkid /dev/vda1
Note
The file system UUID is not fixed; please regularly confirm and update it. For example, after reformatti ng the file system, the UUID of the file system will change.
iii. Execute the following command to open the /etc/fstab file with the VI editor.
vi /etc/fstab
iv. Press i to enter edit mode.
v. Move the cursor to the end of the file, press Enter, and add the following content. Based on the pre vious examples, you would add:
UUID=d489ca1c-xxxx-4536-81cb-ceb2847f9954 / ext4 defaults 0 0
vi. Press Esc, type :wq, and press Enter. Save the settings and exit the editor.
□/etc/fstab: Do not mount other hard disks. After migration, the system may fail to start due to disk l oss.

Uninstall Software
Uninstall drivers and software that may cause conflicts (including VMware tools, Xen tools, Virtualbox GuestAdditions, and some software that comes with its own underlying drivers).
Check Virtio Driver
For details, see Check Virtio Driver in Linux System.
Install cloud-init
For installation details, see Install cloud-init on Linux System.
Check Other Hardware-related Configurations
Hardware changes after migration to the cloud include but may not be limited to:
□The graphics card is replaced with Cirrus VGA.
□The disk is replaced with Virtio Disk, and the device names are vda and vdb.
□The network interface is replaced with Virtio Nic, which only provides eth0 by default.
Find Partitions and Sizes
Execute the following command to view the partition format of the current operating system and det ermine the partition to be copied and its size.
mount
Take the following returned results as an example:
proc on /proc type proc (rw,nosuid,nodev,noexec,relatime)
sys on /sys type sysfs (rw,nosuid,nodev,noexec,relatime)
dev on /dev type devtmpfs (rw,nosuid,relatime,size=4080220k,nr_inodes=1020055,mode=755)
run on /run type tmpfs (rw,nosuid,nodev,relatime,mode=755)
/dev/sda1 on / type ext4 (rw,relatime,data=ordered)
securityfs on /sys/kernel/security type securityfs (rw,nosuid,nodev,noexec,relatime)
tmpfs on /dev/shm type tmpfs (rw,nosuid,nodev)
devpts on /dev/pts type devpts (rw,nosuid,noexec,relatime,gid=5,mode=620,ptmxmode=000)
tmpfs on /sys/fs/cgroup type tmpfs (ro,nosuid,nodev,noexec,mode=755)
cgroup on /sys/fs/cgroup/unified type cgroup2 (rw,nosuid,nodev,noexec,relatime,nsdelegate)
cgroup on /sys/fs/cgroup/systemd type cgroup (rw,nosuid,nodev,noexec,relatime,xattr,name=system d)
pstore on /sys/fs/pstore type pstore (rw,nosuid,nodev,noexec,relatime)
cgroup on /sys/fs/cgroup/cpu,cpuacct type cgroup (rw,nosuid,nodev,noexec,relatime,cpu,cpuacct)
cgroup on /sys/fs/cgroup/cpuset type cgroup (rw,nosuid,nodev,noexec,relatime,cpuset)
cgroup on /sys/fs/cgroup/rdma type cgroup (rw,nosuid,nodev,noexec,relatime,rdma)
cgroup on /sys/fs/cgroup/blkio type cgroup (rw,nosuid,nodev,noexec,relatime,blkio)

cgroup on /sys/fs/cgroup/hugetlb type cgroup (rw,nosuid,nodev,noexec,relatime,hugetlb)
cgroup on /sys/fs/cgroup/memory type cgroup (rw,nosuid,nodev,noexec,relatime,memory)
cgroup on /sys/fs/cgroup/devices type cgroup (rw,nosuid,nodev,noexec,relatime,devices)
cgroup on /sys/fs/cgroup/pids type cgroup (rw,nosuid,nodev,noexec,relatime,pids)
cgroup on /sys/fs/cgroup/freezer type cgroup (rw,nosuid,nodev,noexec,relatime,freezer)
cgroup on /sys/fs/cgroup/net_cls,net_prio type cgroup (rw,nosuid,nodev,noexec,relatime,net_cls,net_prio)
cgroup on /sys/fs/cgroup/perf_event type cgroup (rw,nosuid,nodev,noexec,relatime,perf_event)
systemd-1 on /home/libin/work_doc type autofs (rw,relatime,fd=33,pgrp=1,timeout=0,minproto=5,maxproto=5,direct,pipe_ino=12692)
systemd-1 on /proc/sys/fs/binfmt_misc type autofs (rw,relatime,fd=39,pgrp=1,timeout=0,minproto=5,maxproto=5,direct,pipe_ino=12709)
debugfs on /sys/kernel/debug type debugfs (rw,relatime)
mqueue on /dev/mqueue type mqueue (rw,relatime)
hugetlbfs on /dev/hugepages type hugetlbfs (rw,relatime,pagesize=2M)
tmpfs on /tmp type tmpfs (rw,nosuid,nodev)
configfs on /sys/kernel/config type configfs (rw,relatime)
tmpfs on /run/user/1000 type tmpfs (rw,nosuid,nodev,relatime,size=817176k,mode=700,uid=1000,gid=100)
gvfsd-fuse on /run/user/1000/gvfs type fuse.gvfsd-fuse (rw,nosuid,nodev,relatime,user_id=1000,group_id=100)
It can be seen that the root partition is in /dev/sda1, while there is no independent partition in /boot and /home. And sda1 contains the boot partition, lacks mbr, so we just need to copy the entire sda.
□The exported image must contain at least the root partition and mbr. If the exported image lacks mbr, it will not be able to start.
□In the current operating system, if /boot and /home are independent partitions, the exported image also needs to contain these two independent partitions.
Export Image
Select different ways to export the image based on actual needs.
□Use Command to Export Image
Note
Using commands to manually export images has a high risk (for example, when IO is busy, the file system metadata may become disordered). It is recommended that you check image is complete after exporting it.
You can export the image by executing the following command:
□Use qemu-img command
□Execute the following command to install the required packages. This document uses Debian as an example; package names may differ among various Linux distributions, so please adjust based on your specific situation. For example, in CentOS, the package name is qemu-img.
apt-get install qemu-utils
□Execute the following command to export /dev/sda to /mnt/sdb/test.qcow2.
sudo qemu-img convert -f raw -O qcow2 /dev/sda /mnt/sdb/test.qcow2
If you need to convert to other formats, you can modify the parameter values of -O. The following parameter values can be modified: Among them, /mnt/sdb indicates a new mounted disk or other network storage.

| Parameter Value | Description |

| -------- | -------- |
qcow2	qcow2 format
vpc	vhd format
vmdk	vmdk format
raw	No format

●Use dd command
For example, execute the following command to export the image in raw format.
sudo dd if=/dev/sda of=/mnt/sdb/test.imag bs=1K count=$count
Among them, count parameter is the number of partitions to be copied. You can use fdisk command to find out the value. If you need to copy the entire disk, count parameter can be ignored.
For example, execute the following command to view the number of partitions in /dev/sda.
fdisk -lu /dev/sda
The result resembles the following is returned:
Disk /dev/sda: 1495.0 GB, 1494996746240 bytes
255 heads, 63 sectors/track, 181756 cylinders, total 2919915520 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 4096 bytes
I/O size (minimum/optimal): 4096 bytes / 4096 bytes
Disk identifier: 0x0008f290

Device Boot Start End Blocks Id System
/dev/sda1 * 2048 41945087 20971520 83 Linux
/dev/sda2 41945088 46123007 2088960 82 Linux swap / Solaris
/dev/sda3 46123008 88066047 20971520 83 Linux
/dev/sda4 88066048 2919910139 1415922046 8e Linux LVM
The returned result of the fdisk command shows that the end position of sda1 is 41945087 * 512 byt es, and the count is set to 20481M.
Note
The image exported by the dd command is in raw format. It is recommended to convert the image to qcow2, vhd, or other image format.
Convert Image Format (Optional)
Refer to Convert Image for instructions on how to use qemu-img to convert the image file to a supp orted format.
Check Image
Note
If you create an image without stopping the service or for other reasons, the image file system may b e incorrect. Therefore, it is recommended that you check whether the image is correct after creating i t.
When the image format is consistent with the format supported by the current platform, you can dire ctly open the image to check the file system. For example, the Windows platform can directly attach an image in vhd format, the Linux platform can use qemu-nbd to open an image in qcow2 format, an d the Xen platform can directly enable a vhd file.
1.Sequentially execute the following commands to check whether the nbd module already exists.
modprobe nbd
lsmod | grep nbd
If the returned results are as shown below, then it indicates that the nbd module is present. If the ret urned results are empty, please verify whether the kernel compilation option CONFIG_BLK_DEV_NBD

is enabled. If it's disabled, either switch the system or enable the CONFIG_BLK_DEV_NBD compilation option before recompiling the kernel.

2.Sequentially execute the following commands to check the image.
qemu-nbd -c /dev/nbd0 xxxx.qcow2
mount /dev/nbd0p1 /mnt
After executing the qemu-nbd command, /dev/nbd0 maps the contents of xxx.qcow2. The device /dev/nbd0p1 represents the first partition of this virtual disk; if nbd0p1 does not exist or cannot be mounted successfully, it is likely due to an issue with the image.
Additionally, before uploading the image, you may also start a cloud server to test whether the image file can be used.

Best Practices

Best Practices for CVM

Last Updated At: 2025-08-11 14:49:05

This document helps users maximize the secure and reliable use of Cloud Virtual Machine. Security and Network ●Restrict access: Restrict access by using a firewall (Security Group) to allow trusted address access instances, configuring the strictest rules in the security group. For example, restrict port access, IP address access, etc. ●Security level: Create different security group rules and apply them to instance groups with different security levels to ensure that instances running important businesses cannot be easily accessed from outside. ●Network logical isolation: Select the Virtual Private Cloud for logical partitioning. ●Account permission management: When multiple different accounts are required to control the same set of cloud resources, users can use the Policy Mechanism to control their access permission to cloud resources. ●Secure login: Try to use SSH Token to log in to the user's Linux instance. If you use the Password Log-into log in to the instance, you need to change the password periodically. Storage ●Hardware storage: For data with extremely high reliability requirements, please use Cloud Platform Cloud Block Storage to ensure the persistent storage reliability of data. Try not to select Local Disk. For more information, please see Cloud Block Storage Product Documentation. ● Database: For databases with frequent access and unstable capacity, cloud database services offered by the cloud platform can be utilized. Backup and Restoration ●Backup instance in the same region: You can use Custom Image and Cloud Block Storage Snapshot to back up your instance and business data. For details, see Cloud Block Storage Snapshot and Create Custom Image. ●Cross-region backup instance: You can use Copy Image to copy and back up instances across regions. ●Mask instance failure: You can use Elastic IP to perform domain name mapping to ensure that the service IP can be quickly redirected to another Cloud Virtual Machine instance when the server is unavailable, thereby masking instance failures. Monitoring and Alarms ●Monitor and Event Response: Check monitoring data regularly and set appropriate alarms. For more information, please see Cloud Monitor Product Documentation. ●Burst request processing: Use Auto Scaling to ensure the stability of Cloud Virtual Machine during service peaks and automatically replace unhealthy instances.

Cloud Server Instance Selection Best Practices

Last Updated At: 2025-08-14 09:25:05

This article introduces how to select appropriate server instances based on their features, common business scenarios, precautions, and best practices. Its aim is to guide you in choosing cloud servers suitable for your practical business needs. Region and Availability Zones Region A region defines the geographical location of your cloud computing resources, directly impacting network access conditions for you and your customers to these resources. Availability Zone An availability zone (Zone) is one or more zones within a single region. The types of cloud server instances sold may vary between different zones in the same region. Additionally, there might be differences in network latency when accessing resources across distinct zones. For more information about regions and availability zones, please refer to Regions and Availability Zones. Instance Types The cloud platform offers various instance types, each containing multiple specifications. These can be categorized by architecture into x86 computing, ARM computing, bare-metal computing, heterogeneous computing (GPU), batch computing, etc., or by features such as standard type, memory type, big data type, etc. This article categorizes them based on their feature capabilities, detailed information follows: Standard Type Standard-type instances have balanced performance parameters suitable for most routine business operations like web sites and middleware. Key series under this category include: • S Series: Intel-based core. • Storage Optimized S5se Series: Based on the latest virtualization technology SPDK, specifically optimized for storage protocol stacks, enhancing cloud disk capabilities significantly. Ideal for IO-intensive applications like large databases and NoSQL databases. • Network Optimized SN3ne Series: Capable of up to 6 million packets per second (PPS) over internal networks, an improvement of nearly eight times compared to Standard S3 instances. Can support up to 25Gbps internal bandwidth, which is 2.5 times higher than that of Standard S3 instances. Suitable for high packet traffic scenarios such as video comments, live streaming, gaming, etc. Memory Type M-series Memory Type instances are characterized by large memory capacities, featuring a CPU-to-memory ratio of 1:8 and offering the lowest price per unit of memory. They are primarily applicable to applications requiring extensive memory operations, searches, and calculations such as high-performance databases and distributed memory caching systems like MySQL and Redis. Big Data D-series Big Data Type instances come equipped with massive storage resources and boast high throughput capabilities, making them ideal for Hadoop distributed computing, massive log processing, distributed file systems, and large data warehouse applications that require substantial I/O. Note Local hard disks serve as data disks for D-series big data machine instances, posing risks of data loss (for instance, during host failure). If your application lacks a robust data reliability architecture, we strongly recommend using instances where you can opt for cloud disks as data disks instead. Heterogeneous Computing Heterogeneous computing instances feature GPUs and other heterogeneous hardware, providing real-time high-speed parallel and

floating-point computation capabilities. They are suited for high-performance applications including deep learning, scientific computing, video encoding and decoding, and graphic workstations. Common Business Scenario Selection Recommendations

Business Scenario	Common Software	Scenario Introduction	Recommended Model
Web Service	Nginx Apache	Web services typically include personal websites, blogs, mini-programs, and large e-commerce sites, etc., requiring a balanced demand for computing, storage, memory resources, recommended standard instance configuration meeting business needs.	Standard Series S
Middleware	Kafka MQ	Message queueing businesses require relatively balanced computing and memory resources, recommend standard models equipped with cloud disks for storage.	Standard Series S
Database	MySQL	Databases have very high requirements for IO performance, recommend using SSD cloud disks and local disks (attention needed for data backup when using local disk models due to risk of data loss).	High IO Series IT Memory Type M Series
Caching	Redis Memcache	Cache-based businesses require higher memory but lower computing power, recommend high-memory ratio memory type instances.	Memory Type M Series
Big Data	Hadoop ES	Big data businesses require massive storage and certain IO throughput, recommend dedicated Big Data Series D (attention needed for data backup when using local disk models due to risk of data loss).	Big Data Series D
High Performance Computing	StarCCM WRF-Chem	High-performance computing businesses require extreme single-machine computing power and efficient multi-machine expansion, recommended high-speed RDMA network-equipped high-performance computing clusters or compute type instance families.	High Performance Computing Cluster

Business Scenario	Common Software	Scenario Introduction	Recommended Model
Virtualization	Kvm OpenStack	Virtualization applications require cloud servers to possess nested virtualization capabilities without introducing additional performance overhead, maintaining consistency with traditional physical machine virtualization capabilities. Recommend bare metal cloud server products.	High Performance Computing Cluster Bare Metal Cloud Server
AI Computing	TensorFlow CUDA	AI computing businesses require parallel processing capability, having clear demands for GPU computing power and VRAM.	GPU Compute Type High Performance Computing Cluster

How to build a website?

Last Updated At: 2025-08-14 09:25:05

After you have successfully applied for a cloud server, you can set up your own website or forum on the server you have applied for. Ways to build The cloud platform provides various types of website construction tutorials for mainstream website systems. The ways to build can be divided into image deployment and manual building, each with its respective characteristics:

Comparison item	Image Deployment	Manual Building
Building method	Choose to directly install and deploy through system images from the cloud platform's cloud market.	Install required software manually, which allows customization.
Characteristics	Accompanying software versions are relatively fixed.	Accompanying versions can be flexibly chosen.
Required time	Shorter, one-click deployment.	Longer, requires self-installation of related software.
Difficulty level	Relatively simpler.	Requires certain knowledge about software version compatibility and installation methods.

Build a website You may start building different system personal websites based on your actual needs:

Uploading Local Files to the Cloud Server Uploading files via FTP from a Linux system to the cloud server

Last Updated At: 2025-08-14 09:36:55

Overview This article explains how to use the FTP service on a local machine running the Linux system to upload files from the local machine to the cloud server. Preconditions FTP service has been established on the cloud server. • If your cloud server runs a Linux operating system, consult ”Setting up FTP Service on a Linux Cloud Server” for detailed procedures. • If your cloud server runs a Windows operating system, consult ”Setting up FTP Service on a Windows Cloud Server” for detailed procedures. Directions Connecting to the Cloud Server 1.Execute the following command to install FTP. Note If the local machine running the Linux system already has FTP installed, skip this step and move on to the next one. yum –y install ftp 2.Execute the following command to connect to the cloud server from the local machine, and according to the prompt, enter the username and password for the FTP service. ftp [IP address of the cloud server] Entering the following interface signifies a successful connection.

Uploading Files Execute the following command to upload local files to the cloud server. put [local-file] [remote-file] For example, uploading the local file /home/1.txt to the cloud server. put /home/1.txt 1.txt Downloading Files Execute the following command to download files from the cloud server to the local machine. get [remote-file] [local-file] For example, downloading the A.txt file from the cloud server to the /home directory on the local machine. get A.txt /home/A.txt

Accessing Cloud Object Storage via Intranet from the Cloud Serve

Last Updated At: 2025-08-11 14:57:50

This article introduces the access methods when a CVM (Cloud Virtual Machine) accesses COS (Cloud Object Storage), the judgment method for intranet access, and provides connectivity testing examples. You can refer to this document to gain a deeper understanding of information regarding CVM access to COS. Explanations of Access Methods If you have deployed services inside the cloud platform requiring access to COS, there are distinctions in access methods depending on the region: • Same-region access: Access within the same region range will automatically point to the intranet address, meaning automatic use of the intranet connection. • Cross-region access: Currently does not support intranet access, defaults to resolving to the Internet address. Judgment Method for Intranet Access You can go through this procedure to test whether CVM accesses COS via the intranet: Taking CVM accessing COS as an example, to determine if it uses the intranet to access COS, you can perform the nslookup command on the CVM to resolve the COS domain name. If it returns an intranet IP, it indicates that there is an intranet connection between CVM and COS; otherwise, it is external network access. 1.Obtain the storage bucket access domain name and note down this address. See ”Overview of Storage Buckets” for specifics. 2.Log in to the instance and run the nslookup command. Suppose examplebucket-1250000000.cos.ap-guangzhou.myqcloud.com is the target storage bucket address, then execute the following command: nslookup examplebucket-1250000000.cos.ap-guangzhou.myqcloud.com The result shows that CVM accesses COS via the intranet; the two IP addresses (10.148.214.13 and 10.148.214.14) signify an intranet connection. Note Intranet IP addresses usually resemble 10..., whereas VPC networks often adopt formats like 172.16..*, both fall under the category of intranet.

Testing Connectivity Offers examples for accessing COS over the public internet, same-region CVM (Basic Network) accessing COS, and same-region CVM (VPC Network) accessing COS; for details, see ”Testing Connectivity”.

Operations Guide
Initialize Data Disk
Initialize Data Disk

Last Updated At: 2025-08-14 09:36:55

Initialize Data Disk (Windows Cloud Virtual Machine) Overview After Cloud Virtual Machine is purchased or reinstalled, you need to partition and format the data disk. This document describes how to perform initialization operations such as partitioning and formatting data disks on a Windows Cloud Virtual Machine. Notes ●Formatting the data disk will clear all data. Ensure that there is no data in the data disk or that important data has been backed up. ●To avoid service exceptions, ensure that the Cloud Virtual Machine has stopped external services before formatting. Directions Select the appropriate operation guide according to the disk capacity: ●If the disk capacity is less than 2TB, Initialize Cloud Block Storage (Windows). ●If the disk capacity is greater than or equal to 2 TB, Initialize Cloud Block Storage (Windows). Initialize Data Disk (Linux Cloud Virtual Machine) Overview This document describes how to initialize data disks, such as formatting, partitioning, and creating file systems on a Linux Cloud Virtual Machine. Notes ●Before you format, ensure that there is no data in the data disk or that important data has been backed up. After formatting, all data in the data disk will be cleared. ●To avoid service exceptions, ensure that the Cloud Virtual Machine has stopped external services before you format. Directions Select the appropriate operation guide according to the disk capacity: ●If the disk capacity is less than 2TB, Initialize Cloud Block Storage (Linux). ●If the disk capacity is greater than or equal to 2TB, Initialize Cloud Block Storage (Linux).

Environment Configuration

Environment Configuration

Last Updated At: 2025-08-14 09:36:55

Installing ACPI Power Management Overview In x86 machines, there are two power management methods: APM (Advanced Power Management) and ACPI (Advanced Configuration and Power Interface). ACPI is a power management standard developed jointly by Intel, Microsoft, and Toshiba, providing a more flexible interface for managing computers and devices, whereas APM is an older power management standard. Linux supports both APM and ACPI, but these two standards cannot operate simultaneously. By default, Linux runs ACPI. At the same time, Cloud Platform it is recommended to use ACPI power management. Failure to install ACPI management programs on Linux systems can result in unsuccessful soft shutdowns. This document outlines the steps to check the ACPI installation status and perform the installation. Installation Instructions For CoreOS systems, ACPI installation is not required. Directions 1.Execute the following command to check if ACPI is installed. ps –ef|grep –w ”acpid”|grep –v ”grep” –If the process does not exist, it indicates that ACPI is not installed. Proceed to the next step. –If the process exists, it indicates that ACPI is installed, and the task is complete. 2.Based on your operating system type, execute the corresponding command to install ACPI. –On Ubuntu / Debian, execute the following command: sudo apt–get install acpid –For Redhat / CentOS systems, execute the following command: yum install acpid –For SUSE systems, execute the following command: in apcid

Software Installation

Software Installation

Last Updated At: 2025-08-14 09:48:46

Installing Software on Ubuntu via Apt-get Overview To improve user efficiency of software installation on Cloud Virtual Machine and reduce the cost of downloading and installing software, Cloud Platform provides the Apt-get repository. In the Ubuntu environment, users can quickly install software through Apt-get. For the Apt-get repository, no additional software sources need to be added; software packages can be installed directly. Prerequisites You have logged in to the Cloud Virtual Machine running Ubuntu operating system. Directions Note: The following steps demonstrate how to install Nginx. View available software Execute the following command to view available software. sudo apt-cache search all Software Installation Execute the following command to install Nginx. sudo apt-get install nginx View installed software information Execute different commands based on your actual needs to view information about installed software. ●To view the directory where the software package is located and all the files in that package, execute the following commands. sudo dpkg -L software-name ●Run the following command to view the version information of the software package. sudo dpkg -l software-name Installing Software on CentOS via YUM Overview To improve user efficiency of software installation on Cloud Virtual Machines and reduce the cost of downloading and installing software, Cloud Platform provides a YUM repository. In CentOS environment, users can quickly install the software using the yum command. For YUM repository, users do not need to add additional repositories and can directly install software packages. Directions Software Installation Log in to the Cloud Virtual Machine using the root account. Execute the following command to install the software. CentOS 8 and later versions 1.Execute the following command to install the software. dnf install software-name During the software installation process, the system will automatically search for relevant software packages and dependencies, and prompt the user to confirm whether the found packages are suitable. 2.After confirming that the software package is correct, type y and press Enter to start installing the software. CentOS 7 and lower versions 1.Execute the following command to install the software. yum install software-name 2.After confirming that the software package is correct, type y and press Enter to start installing the software. Viewing Installed Software Information After the software installation is complete, you can execute different commands based on actual needs to view information. ●Run the following command to view the specific installation directory of the software package. pm -ql software-name ●Run the following command to view the version information of the software package. rpm -q Installing and Configuring Chronyd Service on CentOS 8 Overview The native CentOS 8 does not support installing the ntp service, leading to potential issues with inaccurate system time. Therefore, it's necessary to adjust the time service using chronyd. This article introduces how to install and configure the chronyd time service on a cloud platform server running CentOS 8 operating system. Directions Installation and Configuration of Chronyd Service

1.Log in to the cloud server instance; for details, see Logging in to a Linux Instance Using Standard Method (Recommended). 2.Execute the following command to install the chronyd service. yum –y install chrony 3.Execute the following command to modify the configuration file chrony.conf. vim /etc/chrony.conf 4.Press i to enter edit mode, then under the line commented out by #log measurements statistics tracking, add the following lines. server time1.yun.com iburst server time2.yun.com iburst server time3.yun.com iburst server time4.yun.com iburst server time5.yun.com iburst 5.Press Esc and type :wq to save and exit the editing mode. 6.Sequentially execute the following commands to set the chronyd service to start automatically upon boot and restart the service. systemctl restart chronyd systemctl enable chronyd Check Service Configuration 1.Execute the following command to check whether the time is synchronized. date 2.Execute the following command to view the status of time synchronization sources. chronyc sourcestats —v Appendix Common Commands

Command	Description
chronyc sources –v	View time synchronization sources.
chronyc sourcestats –v	View the status of time synchronization sources.
timedatectl set–local–rtc 1	Set hardware time, which is by default in UTC.
timedatectl set–ntp yes	Enable NTP time synchronization.
chronyc tracking	Calibrate the time server.
chronyc –a makestep	Force synchronization of system clock.

Setting Custom Data
Setting Custom Data (Linux Cloud Virtual Machine)

Last Updated At: 2025-08-14 10:25:12

Overview When creating Cloud Virtual Machine, you can configure the instance by specifying custom data. When Cloud Virtual Machine starts for the first time, the custom data is passed to the cloud server as text and executed. If you purchase multiple Cloud Virtual Machines at once, the custom data will be executed on all cloud servers upon their first startup. This document uses an example of passing a PowerShell script during the first startup of a Linux Cloud Virtual Machine. Notes ●Supported Linux operating systems for custom data include: -64-bit OSs: CentOS 6.8 and later, Ubuntu Server 14.04.1 LTS and later, and openSUSE 42.3 x86. -32-bit Operating Systems: CentOS 6.8 32-bit and above ●Commands are executed by passing text only during the initial startup of the Cloud Virtual Machine. ●Note: The text must be Base64 encoded. Please encode the text in a Linux environment to avoid format incompatibility. ●Use the root account to execute the user data input text. Do not use the "sudo" command in the script. Any files you create will be owned by root. If you need non-root users to have file access, modify the permissions in the script. ●Executing tasks specified in custom data during startup will increase the server's startup time. Please wait a few minutes and test if the task has been successfully executed. ●In this example, the Shell script must start with #! character and the path to the interpreter that will read the script Directions Writing the Shell Scripts Execute the following command to create a Shell script file named 1."script_text.sh". vi script_text.sh 2.Press i Switch to edit mode, refer to the following content, write and save the "script_text.sh" script file. #!/bin/bash

Encoding Script Files Using Base64 1.Execute the following command to Base64 encode the script_text.sh script file.

Base64 encode the script.

base64 script_text.sh The following information will be returned:

The encoded result will be:

2.Run the following command to verify the result of Base64 encoding script.

Perform Base64 decoding on the returned result to verify whether it is the command to be executed

Transmitting through API When you create a CVM through an API, you can assign the encoded result returned from Encoding Script Files Using Base64 to the UserData parameter of the RunInstances API to pass the text.\n For example, create a request parameter for a Cloud Virtual Machine with a UserData parameter, as shown below:

View Execution Logs After successfully creating the server, you can execute the following command to view the script execution logs:

System Related

System Related

Last Updated At: 2025-08-14 10:25:12

Windows Recovery Mode What is Windows Recovery Mode? Windows System Recovery Mode (Recovery) indicates that Windows uses the automatic repair feature. When Windows detects certain system issues and believes that continued use will cause damage to the system, it will prevent Windows from starting and enter the System Recovery Options interface, providing a state for users to repair, backup, or restore the system. The System Recovery Options include several tools, such as Startup Repair, System Restore and Windows Memory Diagnostics. These tools can be used to fix problems, backup data, and perform system recovery. When the user cannot log in to the CVM remotely and the following status appears after logging in to the CVM through the console, it indicates that the Windows CVM has entered the recovery mode. Reasons for Entering Recovery Mode Common reasons for entering the recovery mode are as follows: The power is forcefully turned off while Windows is running or shutting down. This includes forced shutdowns performed in the console. Improper shutdown may cause the system to lose some critical data and thus enter recovery mode. The power is cut off during Windows update. Critical data is lost during the update and the system enters the recovery mode. The system is damaged by a Trojan or virus. There is a bug in Windows core services. A risk is detected through Windows self-check. The system loses critical data or is damaged. Users may accidentally damage system files, causing the system to enter the recovery mode. Precautions Cloud Platform recommends that users take the following precautions: When shutting down, open the console to observe the Windows shutdown process. Cloud Platform soft shutdown has a timeout mechanism. After a soft shutdown is performed, if the system does not shut down within the scheduled time, a prompt of failure will pop up. If the shutdown is slow or Windows update occurs during shutdown, just wait for the CVM to shut down and do not force shutdown. It is recommended to refer to Several Scenes of Shutdown Failure. Check whether there are any abnormal programs or processes such as Trojans or viruses in the system. Check whether the system management and anti-virus software are running normally. Keep Windows update packages up-to-date in a timely manner, especially some important updates and security updates. Regularly check the system event log to verify if there are any errors in the core services. Solution After Windows enters the recovery mode, users can try to continue startup and running, or perform automatic recovery. Minor problems can be repaired by Windows itself. Execute the following steps: 1.Log in to the CVM from the console. 2.Recovery mode interface appears, click Next. 3.System Recovery options appear, click Next, using the default plan. 4.click Restart, and quickly press the keyboard F8. 5.Select Normal Boot of Windows. 6.If Windows cannot start, reinstall the system in the console. For details, refer to Using the Console to Reinstall the System. Windows System Update Updating via Public Network You can install patches using the system's Windows Update service. Follow the steps below: 1.click > Control Panel > Windows Update,

open the Windows Update window. 2.Click Check for Updates, and wait for the check to complete. 3.After the check is complete, click Windows Update to see an important updates available or n optional updates available. 4.In the pop-up Select Updates to Install window, select the updates you want to install, click Install.

After the update is installed, if the system prompts that the system needs to be restarted, please restart the CVM in time. Note: When you restart the CVM after applying updates, you need to log in and monitor the CVM through VNC. If the system shows prompts such as ”Updating, do not turn off the power” or ”Configuration not complete”, then do not perform a hard shutdown. A hard shutdown may damage your CVM. Updating via private network If your CVM cannot connect to the public network, you can use the Cloud Platform private network patch server to install updates. The Cloud Platform Windows patch server includes most common updates for Windows but does not include hardware driver packages or some less frequently used server updates. Some less common services might not have updates available on the Cloud Platform private network patch server. How to use the Cloud Platform private network patch server is as follows? 1.Log in to the Windows CVM. 2.Use IE browser to access and download the Cloud Platform private network setting tool (wusin.bat). 3.Save the ’wusin.bat’ setting tool to the C: drive. Use the administrator Cloud Platform Command Line Interface (CMD) to open wusin.bat. Note: If you run the ’wusin.bat’ tool directly via IE, the console window will automatically close, and you won’t be able to observe the output information. If you no longer need to use the Cloud Platform private network Windows patch server, you can download the cleanup tool wusout.bat to clean up. The method is as follows: 4. 1.Log in to the Windows CVM. Use IE browser to access and download the Cloud Platform intranet cleanup tool (wusout.bat). 2.Save the wusout.bat setup tool to the C: drive. Use the Administrator Cloud Platform Command Line Interface (CMD) to open wusout.bat. As shown below:

Note: If you execute the wusout.bat tool directly through IE, the console window will be automatically closed and the output information cannot be observed. Windows Server System Activation The Cloud Virtual Machine uses KMS to authorize Windows Servers Licensing. Only Windows Server 2008 and Windows Server 2012 require this activation method. Activation Prerequisites 1.In Windows, the SPP Notification Service is used for activation-related services and must be kept running. 2.Some optimization software may disable the ability to modify the execution permissions of service-related executables, such as sppsvc.exe. If the execution permissions of sppsvc.exe are modified, the service may not run properly. Before attempting to activate a Windows Cloud Virtual Machine, ensure that this service and other basic functions on Windows are operating normally. Automatic activation Cloud Platform provides a script to simplify the manual activation steps for Windows Servers. Manual activation 1.Log in to the Windows CVM. 2.On the operating system interface, click Start then Run, enter cmd.exe to open the command console window. 3.In the console window, execute the following commands in sequence

cscript /nologo %windir%/system32/slmgr.vbs -skms kms.cce.yun.ccb.com:1688 cscript /nologo %windir%/system32/slmgr.vbs —ato

Execute the above commands in the console window to complete manual activation Note: on certain systems, if there's an issue with the system clock, errors may occur during manual activation. In this case, you need to synchronize the system clock first. The method to sync the clock is: Enter the following command in the console window:

w32tm /config /syncfromflags:manual /manualpeerlist:"ntpupdate.yfm4-v6-iaas.tcecloud.fsphere.cn/"
w32tm /resync

Modifying SID Operation Instructions Note: This instruction applies only to Windows Server 2008 R2 and Windows Server 2012 systems. If there is a bulk modification requirement for SIDs, it can be resolved by creating a custom image (select "perform sysprep to create an image"). Overview Security identifiers (SIDs) are used in Microsoft operating systems to identify computers and users. Due to the same SIDs of CVM instances produced based on the same image, it may cause the problem of being unable to join the domain. If you need to build a Windows domain environment, you need to modify the SIDs to join the domain. Directions 1.Log in to the CVM through VNC. 2.On the operating system interface, right-click click Run, type cmd, execute the following command to save the current network configurations. ipconfig /all 3.In the administrator TCCLI, execute the following command to open the sysprep tool. C:\Windows\System32\Sysprep\sysprep.exe -Set System Cleanup Operation to Enter the System Out-of-Box Experience (OOBE) and check General. -Set Shutdown Options to Restart. 4.Click OK, and the system will automatically restart.After the startup is complete, follow the wizard to complete the configurations (by selecting the language, resetting the password, etc.). 5.On the operating system interface, right-click click Run, type cmd, Run the following command to verify whether the SIDs have been modified. whoami /user 6.According to the network configuration information saved in Step 2, reset the network interface information (such as IP address, gateway address, DNS, etc.). Updating Virtio Network Interface Drivers Overview Cloud Platform CVM Window Server 2008 R2 Enterprise Edition SP1 and Windows Server 2012 R2 optimize the network performance of virtual hardware through installing Virtio network interface drivers. Cloud Platform will continue to improve the network interface drivers to enhance performance and rectify faults. You can view your system version information through following steps: log in to CVM, and right-click on the desktop click Computer > Properties, to open the System window.In System under View Basic Computer Information, you can view the system version information. Methods to Update Virtio Network Interface Drivers Note: The network may be intermittently disconnected during the update process. Before updating, check if it will affect your business. After the update, you will need to restart your computer. 1.Log in to the CVM 2.Download the installation files of VirtIO network interface driver for Window Server 2008 R2 and Windows Server 2012 R2 through the browser in CVM. 3.After downloading, double-click to start the installer, choose the Typical installation mode, click Next.

5.During the installation, if the following pop-up window appears, please select Always install this driver software. Follow the prompts and restart your computer to complete the update. Configuring High-performance Power Management Overview On the Windows Server operating system, high-performance power management needs to be configured to support soft shutdown of instances. Otherwise, the Cloud

Virtual Machine console can only shut down the instance through hard shutdown. This document uses Windows Server 2012 operating system as an example to describe how to configure power management. Note: Modifying power management does not require a computer restart. 1.Log in to the Windows CVM. 2.Access the Cloud Platform private network through IE browser and download the Cloud Platform Power Management Modification and Configuration Tool (power–set–win.bat) to the C: drive. 3.Open power–set–win.bat using the administrator Cloud Platform Command Line Interface (CMD). Run the following command to view the current power management scheme. powercfg –L 4.In the operating system interface, click > Control Panel > System and Security > Power Options, to open the Power Options window, click Change plan settings. In the opened ”Edit Plan Settings” window, modify the idle shutdown time for the monitor and hard drive. System Language Adjustment

The Windows system provided by the cloud platform defaults to the Chinese version. We have added an English language pack to the machine. Users who require this change can modify the language themselves. The specific steps are as follows: 1)Enter the clock, language, and region configuration interface. Click [Start] – [Control Panel] – [Clock, Language, and Region]. 2)Configure the system language environment to English. Click [Region and Language] – [Keyboard and Language] – [Select Display Language], choose English, then click Apply. 3)Log out and log back in; the new language environment will take effect.

The Windows system provided by the cloud platform defaults to the Chinese version. If you changed it to English and want to switch back to Chinese, please follow these steps. 1)Enter the clock, language, and region configuration interface. Click [Win] – [Control Panel] – [Clock, Language, and Region] 2)Configure the system language environment to Chinese. Click [Region and Language] – [Keyboard and Language] – [Choose a display language] – select ”Simplified Chinese”. 3)Log out and log back in; the new language environment will take effect.

Introduction to Common Kernel Parameters for Linux Instances

Last Updated At: 2025-08-14 11:11:37

The cloud platform has already configured some parameters by default in the public Linux images. However, due to the highly personalized configuration capabilities of sysctl, the cloud platform recommends users configure sysctl according to their own business characteristics. You can learn about the specific default optimization configurations for public cloud Linux images and common configurations through this article, and manually optimize based on your business needs. Note • For parameter items where the ”initialization configuration” is ”–”, they all maintain the default configuration of the official image. • Using the command sysctl –w configures temporary effects; writing into /etc/sysctl.conf configures permanent effects. Network Category

Parameter	Description	Initialization Configuration
net.ipv4.tcp_tw_recycle	Enables quick recycling of TIME_WAIT connections. Disabling prevents kernel from checking packet timestamps, whereas enabling performs such checks. Not recommended, could cause packet loss issues under certain conditions. Removed in newer kernels.	0
net.core.somaxconn	Relates to the established state post three–way handshake completion, without an accept queue. High accept queues suggest low server–side efficiency or sudden bursts of new connections. Smaller values might prevent SYN responses due to SOMAXCONN table overflow.	128
net.ipv4.tcp_max_syn_backlog	Upper limit for half–open connections, historically utilized for defending against synflood attacks. Semi–connected states can exceed this limit when tcp_syncookies=1.	–

Parameter	Description	Initialization Configuration
net.ipv4.tcp_syncookies	Enables SYN Cookies for handling situations, capable of preventing some SYN attacks and maintaining connectivity despite SYN wait queue overflow. Uses SHA1 validation upon enabling, possibly increasing CPU usage.	1
net.core.rmem_default net.core.rmem_max net.ipv4.tcp_mem net.ipv4.tcp_rmem	These parameters configure the size of data reception buffers. Overly large configurations can easily result in wastage of memory resources, while too small a setting may lead to packet loss. It's recommended to determine whether your business falls under high-concurrency connections or fewer concurrent high-throughput scenarios for optimized configuration.• The theoretically optimal configuration strategy for rmem_default is the product of bandwidth and RTT, its configuration overrides tcp_rmem, which isn't configured separately.• rmem_max is approximately five times the size of rmem_default.• tcp_mem refers to total TCP memory utilization, usually auto-configured by the OS to 3/32, 1/8, or 3/16 of available CVM memory, and both tcp_mem and rmem_default also determine the maximum number of concurrent connections.	rmem_default =655360 rmem_max =3276800

Parameter	Description	Initialization Configuration
net.ipv4.tcp_keepalive_intvl net.ipv4.tcp_keepalive_probes net.ipv4.tcp_keepalive_time	These parameters relate to TCP KeepAlive, with default settings of 75/9/7200. This means that for any given TCP connection, the kernel will start probing after being idle for 7200 seconds, and if unsuccessful after 9 probes (each lasting 75 seconds), it will send an RST packet. For servers, the default values are relatively large, but they can be adjusted to 30/3/1800 in conjunction with business needs.	–
net.ipv4.ip_local_port_range	Specifies usable port range, adjust as needed.	–
tcp_tw_reuse	Allows reuse of sockets in TIME-WAIT state for new TCP connections. Useful for rapid reconnection but poses risks in NAT networks. Newer kernels offer three options: 0/1/2, with a setting of 2.	–
net.ipv4.ip_forward net.ipv6.conf.all.forwarding	IP forwarding functionality, set to 1 for Docker-related routing scenarios.	0
net.ipv4.conf.default.rp_filter	Reverse path filtering rules for incoming packets, configurable to 0/1/2. Based on RFC3704 recommendations, set to 1 for strict filtering to counteract DDoS attacks and IP spoofing.	–
net.ipv4.conf.default.accept_source_route	Disallows acceptance of source-routed IP packets by default, per CentOS guidelines.	0
net.ipv4.conf.all.promote_secondaries net.ipv4.conf.default.promote_secondaries	Determines promotion of secondary IPs to primary status upon removal of primary IP.	1

Parameter	Description	Initialization Configuration
net.ipv6.neigh.default.gc_thresh3 net.ipv4.neigh.default.gc_thresh3	The limitation on the maximum number of records stored in the ARP cache. Once the number of entries in the cache exceeds the specified value, the garbage collector will immediately run.	4096

Memory Category

Parameter	Description	Initialization Configuration
vm.vfs_cache_pressure	Original value is 100, indicating the intensity of scanning dentries. Using 100 as a baseline, the higher this value, the more inclined the kernel's reclaim algorithm is towards freeing up memory. Many curl-based services often experience filling up all available memory due to accumulated dentries, leading to OOM or kernel bugs. After considering reclaim frequency and performance, choosing a configuration of 250 allows for adjustment as necessary.	250
vm.min_free_kbytes	The value is calculated at startup based on system physical memory (MEM): 4 * sqrt(MEM). Its meaning is the minimum amount of KB memory that should be reserved during system operation, typically provided for kernel threads. There's no need to set this value excessively high. When packet volumes exhibit micro-bursts, there's a chance of breaching vm.min_free_kbytes, causing OOM. It's suggested to configure machines with ample resources to set vm.min_free_kbytes to around 1% of total memory by default.	-
kernel.printk	The print level of the kernel printk function, default configuration is greater than 5.	5 4 1 7

Parameter	Description	Initialization Configuration
kernel.numa_balancing	This parameter indicates that the kernel can autonomously move process data to corresponding NUMA nodes, however, practical application outcomes are suboptimal and accompanied by other performance impacts. Under Redis scenarios, attempting to enable this feature might be beneficial.	0
kernel.shmall kernel.shmmax	• shmmax sets the maximum length for allocating shared memory in bytes at once.• shmall defines the overall maximum length of shared memory that can be allocated, measured in pages.	kernel.shmmax=68719476736 kernel.shmall=4294967296

Process Category

Parameter	Description	Initialization Configuration
fs.file-max fs.nr_open	Controls the maximum number of files all processes and each process can open simultaneously:• file-max is automatically configured by the OS upon startup, approximately 100,000 per GB.• nr_open is a fixed value of 1,048,576, but it limits the maximum number of files opened by the user space. This value is generally not modified; instead, ulimit -n is usually set, corresponding configuration file is /etc/security/limits.conf.	The open files limit for ulimit is 100,001 and fs.nr_open=1048576
kernel.pid_max	Maximum number of processes in the system. In official images, the default is 32,768, which can be adjusted as needed.	-
kernel.core_uses_pid	This setting determines whether the PID is included when generating core dump files.	1
kernel.sysrq	After enabling this parameter, subsequent operations can be performed on /proc/sysrq-trigger.	1

Parameter	Description	Initialization Configuration
kernel.msgmnb kernel.msgmax	Represents the maximum number of bytes in message queues and the capacity of individual largest message queue.	65536
kernel.softlockup_panic	When softlockup_panic is configured, the kernel will panic upon detecting a soft lockup in any process. Combined with kdump settings, a vmcore can be generated to analyze the cause of the soft lockup.	–

IO category

Parameter	Description	Initialization Configuration

Parameter	Description	Initialization Configuration
vm.dirty_background_bytes vm.dirty_background_ratio vm.dirty_bytes vm.dirty_expire_centisecs vm.dirty_ratio vm.dirty_writeback_centisecs	These parameters primarily configure the strategy for writing back IO to the disk: dirty_background_bytes and dirty_bytes, along with dirty_background_ratio and dirty_ratio, correspond respectively to the absolute amount and proportional amount of dirty page thresholds in memory. Typically, ratios are set.□dirty_background_ratio indicates the percentage of system memory that contains dirty pages in the filesystem cache (default 10%) at which point the kernel's flush processes and others are awakened to write back to the disk.□dirty_ratio represents the maximum proportion of dirty pages. When the number of dirty pages reaches this ratio, all dirty data must be submitted to the disk, and all new IO operations will be blocked until the dirty data is written to the disk, which commonly results in IO latency. Initially, the system meets the conditions set by vm.dirty_background_ratio, then triggers asynchronous write-back operations through flushing processes. During this phase, application processes can still perform write operations. However, if the value set by vm.dirty_ratio is reached, the operating system transitions to synchronously handling dirty pages, blocking application processes.□vm.dirty_expire_centisecs denotes the survival time for dirty pages. Flushing processes check whether the data has exceeded this time limit, measured in hundredths of a second.□vm.dirty_writeback_centisecs signifies the wake-up cycle for flushing processes, also measured in hundredths of a second.	–

xx

Troubleshooting

Instance-related Failures

Instance-related Failures

Last Updated At: 2025-08-14 11:11:37

High Bandwidth Usage Prevents Log-in This document introduces the troubleshooting methods and solutions for Linux and Windows CVM when they cannot be remotely connected due to high bandwidth usage. Failure Symptoms □After the CVM console log-in, the bandwidth monitoring data of the CVM prompts that excessively high bandwidth usage is preventing connection to the Cloud Platform server. □Diagnosed excessively high bandwidth usage through self-diagnostic tools. Failure Location and Troubleshooting For Linux Servers After you log in to the Linux CVM via VNC, you need to do the following: Note: The following operations take the CVM running CentOS 7.6 as an example. 1.Run the following command to install the iftop tool (the iftop tool is a utility for monitoring traffic on Linux servers). yum install iftop -y Note: For Ubuntu system, please run apt-get install iftop -y command. 2.Run the following command to install lsof. yum install lsof -y 3.Run the following command to execute iftop. iftop -<=,=>indicates the direction of traffic. -TX means sending traffic. -RX means receiving traffic. -TOTAL indicates total traffic. -Cum indicates the total traffic since iftop was run up to now. -peak indicates the peak flow rate. rates represent the average traffic rate over the past 2s, 10s, and 40s. 4.According to the IP consuming traffic in iftop, execute the following command to track the process connected to the IP. lsof -i | grep IP For example, if the IP consuming traffic is 201.205.141.123, execute the following command: lsof -i | grep 201.205.141.123 According to the following returned results, we know that the bandwidth of this server is mainly consumed by the SSH process. sshd 12145 root 3u IPV4 3294018 0t0 TCP 10.144.90.86:ssh->203.205.141.123:58614(ESTABLISHED) sshd 12179 ubuntu 3u IPV4 3294018 0t0 TCP 10.144.90.86:ssh->203.205.141.123:58614(ESTABLISHED) 5.Check the process that consumes more bandwidth to evaluate whether the process is normal. -If a business process is consuming more bandwidth, you need to analyze whether this increase is caused by changes in access volume and whether space optimization or server configuration upgrades are required. -If an exception process consumes more bandwidth, this increase may be caused by a virus or Trojan horse. You can terminate the process manually or use security software to detect and kill it. You can also back up the data and reinstall the system. Failure to shut down and restart cloud server When performing operations such as shutting down and restarting a cloud server, there is a very small probability of failure to shut down or restart. If you encounter such situations, you can troubleshoot and handle the cloud server as follows. Possible causes • Excessively high CPU or memory utilization rate. • The Linux operating system-based cloud server does not have the ACPI management program installed. • System updates for Windows

operating system–based cloud servers take too long. • When applying for Windows cloud servers for the first time, these servers haven't been initialized. • Some software was installed on the operating system, or it got infected by trojans or viruses, resulting in damage to the system itself, etc. Troubleshooting Check CPU/Memory Usage 1.Check the CPU/memory usage based on the type of the cloud server's operating system. o For Windows cloud servers: Right–click the taskbar in the cloud server and select Task Manager. o For Linux cloud servers: Execute the top command to view the information in the %CPU column and the %MEM column. 2.Terminate processes with excessively high CPU or memory utilization rates according to actual CPU/memory usage. If you still cannot shut down/restart, perform the force shutdown/restart function. Check if the ACPI Management Program is Installed Note This operation targets Linux operating system–based cloud servers. Execute the following commands to check whether the ACPI process exists. ps –ef | grep –w "acpid" | grep –v "grep" If the ACPI process exists, perform the force shutdown/restart function. • If there is no ACPI process, install the ACPI management program. For specific operations, refer to Configuring Linux Power Management. Check if Windows Update is Running Note This operation targets Windows operating system–based cloud servers. In the Windows cloud server OS interface, click Start > Control Panel > Windows Updates to see if any patches or programs are being updated. • When Windows performs certain patch operations, it will do some processing when shutting down the system. At this point, the update may take too long, leading to failure to shut down/restart. It is recommended that you wait until Windows updates are complete before attempting to shut down/restart the cloud server. • If there are no patches or programs being updated, perform the force shutdown/restart function. Check if the Cloud Server Has Been Initialized Note This operation targets Windows operating system–based cloud servers. When applying for Windows cloud servers for the first time, the system distributes images via the Sysprep method, which takes slightly longer during initialization. Before initialization is complete, Windows will ignore shutdown/restart operations, causing failure to shut down/restart. • If the Windows cloud server you applied for is initializing, it is recommended that you wait until the initialization of the Windows cloud server is complete before attempting to shut down/restart the cloud server. • If the cloud server has already been initialized, perform the force shutdown/restart function. Check if the Installed Software is Normal Use inspection tools or antivirus software to check if the software installed on the cloud server is normal or has been infected by trojans, viruses, etc. • If abnormalities are found, it indicates that the system itself might have been damaged, leading to failure to shut down/restart. It is suggested that you uninstall the software, scan using security software, or after backing up data, reinstall the system. • If no abnormalities are found, perform the force shutdown/restart function. Force Shutdown/Restart Function Note The cloud platform provides a force shutdown/restart feature, which can be used when multiple attempts to shut down or restart the cloud server fail. This operation forcefully shuts down or restarts the cloud server, potentially causing loss of data on the cloud server or damage to the file system. 1.Log in to the cloud server console. 2.On the instance management page, select the cloud server to be shut down or restarted and perform different operations according to actual needs. o Shut down the cloud server: Click More > Instance Status > Shut Down. o Restart the cloud server: Click More > Instance Status > Restart. 3.In the pop–up window titled "Shutdown" or

”Restart Instance,” select ”Force Shutdown” or ”Force Restart” and click OK. o Select ”Force Shutdown” o Select ”Force Restart” Unable to Create Network Namespace Problem Description When performing a command to create a new Network Namespace, the command gets stuck and cannot continue. Dmesg message prompt: ”unregister_netdevice: waiting for lo to become free. Usage count = 1” Causes of Problem This problem is a kernel bug. Currently, the following kernel editions have this bug: ●Ubuntu 16.04 x86_64 kernel edition is 4.4.0−−91−generic. ●Ubuntu 16.04 x86_32 kernel edition is 4.4.0−−92−generic. Solution Upgrade the kernel edition to 4.4.0−−98−generic, which has fixed this bug. Processing Procedures 1.Perform the following command to check the current kernel edition. uname −r 2.Perform the following command to check whether the 4.4.0−−98−generic kernel edition is available for upgrade. sudo apt−get update sudo apt−cache search linux−image−4.4.0−98−generic If the following information is displayed, it represents that the edition exists in the source and can be upgraded: linux−image−4.4.0−98−generic − Linux kernel image for version 4.4.0 on 64 bit x86 SMP 3.Perform the following command to install the new edition of the kernel and the corresponding Header package. sudo apt−get install linux−image−4.4.0−98−generic linux−headers−4.4.0−98−generic 4.Perform the following command to restart the system. sudo reboot 5.Perform the following command to enter the system and check the kernel edition. uname −r If the following result is displayed, it represents the edition update is successful: 4.4.0−98−generic Kernel and IO Related Problems Kernel Problem Location and Solution Failure Symptoms Kernel−related failures may cause the machine to be unable to log in or abnormal restart. Possible Reasons Kernel hung_task The hung task mechanism is implemented by the kernel thread khungtaskd, which monitors the process in the TASK_UNINTERRUPTIBLE status. If in the period of kernel.hung_task_timeout_secs (Default 120 seconds), it remains in D status, then the stack information of the hung task process will be printed. If configure kernel.hung_task_panic=1, it will trigger a kernel panic and restart the machine. Kernel soft lockup Soft lockup means that the CPU is occupied by kernel code and cannot perform other processes. The principle of detecting soft lockup is to assign a kernel thread [watchdog/x] which will be timed perform to each CPU. If the thread is within a certain period (the default is 2*kernel.watchdog_thresh, 3.10 kernel kernel.watchdog_thresh and the default value is 10 seconds) is not performed, indicating that a soft lockup has occurred. If kernel.softlockup_panic=1 is configured, it will trigger a kernel panic and restart the machine. Kernel panic The kernel’s abnormal crash causes the machine to restart abnormally. Common kernel panic scenarios are as follows: □The kernel has a hung_task and kernel.hung_task_panic=1 is configured. □The kernel has a soft lockup and kernel.softlockup_panic=1 is configured. □The kernel bug is triggered. Processing Procedures The troubleshooting and processing steps for kernel−related problems are relatively complicated. It is recommended to submit a ticket for further location and processing. Hard Disk Problem Location and Solution Hard disk inode is full Failure phenomenon: When creating a new file, the error message No space left on device is prompted, and the inode space usage is 100% by using the df −i command. Possible causes: File system inode exhaustion. Processing steps: Delete unnecessary files or scale out the hard disk capacity. Hard disk space usage is full Failure phenomenon: When creating a new file, the error message No space left on device is prompted, and the hard disk space usage is 100% by using the df −h

command. Possible causes: Hard disk space exhaustion. Processing steps: Delete unnecessary files or scale out the hard disk capacity. Hard disk read only Failure phenomenon: The file system can only read files but cannot create new files. Possible cause: The file system is damaged. Processing steps: 1. 1.Create a snapshot to back up disk data. See Create a Snapshot for details. 2.Perform the corresponding processing steps according to the hard disk type: □System disk □Data disk It is recommended to directly restart the instance; for details, see Restarting an Instance. Hard disk %util high Failure phenomenon: The instance is stuck, and logging in using SSH or VNC is slow or unresponsive. Possible causes: High IO causes hard disk %util to reach 100%. Processing steps: Check whether the high IO is reasonable, and evaluate whether to reduce IO read and write or replace the hard disk with a higher performance. Missing Soft Links in System Bin or Lib Symptom Description During command execution or system startup, errors occur indicating that commands or lib libraries cannot be found. Possible Causes In systems such as CentOS 7, CentOS 8, and Ubuntu 20, directories such as bin, sbin, lib, and lib64 are soft links. As shown below: lrwxrwxrwx 1 root 7 Jun 19 2018 bin –> usr/bin lrwxrwxrwx 1 root 7 Jun 19 2018 lib –> usr/lib lrwxrwxrwx 1 root 9 Jun 19 2018 lib64 –> usr/lib64 lrwxrwxrwx 1 root 8 Jun 19 2018 sbin –> usr/sbin If these soft links are deleted, errors will occur during command execution or system startup. Resolution Ideas Refer to the handling steps to check and create required soft links. Directions 1.Refer to Using Rescue Mode to enter rescue mode. 2.Execute the mount and chroot commands mentioned therein. When executing the chroot command: o If there's an error, execute cd /mnt/vm1. o If there's no error, execute cd /. 3.Execute the following command to check if the corresponding soft link exists. ls –al / | grep –E "lib|bin" o If yes, please contact us for assistance through Online Support. o If no, then please execute the following commands as needed to create the corresponding soft link. o ln –s usr/lib64 lib64 o ln –s usr/sbin o ln –s usr/bin bin ln –s usr/lib lib 4.Execute the following command to verify the soft links. chroot /mnt/vm1 /bin/bash If there are no error messages, it means the soft links have been successfully repaired. 5.Refer to Using Rescue Mode to exit rescue mode and boot the system. Error creating files due to "no space left on device" Symptom Description When creating new files in a Linux cloud server, an error message "no space left on device" appears. Possible Causes • Hard disk space is full. • File system inode is full. • Inconsistency between df and du. o A file has been deleted but there are still processes holding onto the corresponding file descriptor, preventing the hard disk space from being released. o Nested mounts. For example, the /data directory on the system disk occupies a large amount of space, and /data is also mounted to other data disks, leading to inconsistencies between df and du on the system disk. Resolution Ideas Refer to Troubleshooting Methods to identify and resolve issues. Troubleshooting Methods Resolving Hard Disk Space Full Issues 1.Log in to the cloud server; for details, see Logging into a Linux Instance using Standard Login Method. 2.Execute the following command to view hard disk usage. df –h 3.Identify the mount point with higher hard disk usage and enter the mount point with the following command. cd corresponding_mount_point For example, to cd to the system disk mount point, execute cd /. 4.Execute the following command to find directories occupying significant space. du –x ––max–depth=1 | sort –n Based on the capacity of the largest occupied directory identified, proceed as follows: o If the directory size is much lower than the total hard disk space, continue troubleshooting per

the Inconsistent df du Problem Resolution steps. o If the directory size is considerable, locate larger files within the identified directory in step 2. Assess whether they can be deleted considering business requirements. If deletion isn't possible, expand cloud disk storage to increase hard disk capacity. Resolving File System Inode Full Issues 1.Log in to the cloud server; for details, see Logging into a Linux Instance using Standard Login Method. 2.Execute the following command to view hard disk usage. df –i 3.Identify the mount point with higher hard disk usage and enter the mount point with the following command. cd corresponding_mount_point For example, to cd to the system disk mount point, execute cd /. 4.Execute the following command to find the directory containing the most files and address the issue. Please note that this command can be time–consuming, so be patient. find / –type f | awk –F / –v OFS=/ '{NF="";dir[NF="";dir[0]++}END{for(i in dir)print dir[i]" "i}' | sort –k1 –nr | head Resolving Inconsistent df du Problems Addressing Process Holding onto File Descriptors Execute the following command to view processes holding onto files. lsof | grep delete Based on the returned results, proceed as follows: • Kill the corresponding process. • Restart services. • If many processes hold onto file descriptors, consider rebooting the server. Addressing Nested Mounts Issue 1.Execute the mount command to mount the disk occupying large space to /mnt. For example: mount /dev/vda1 /mnt 2.Execute the following command to enter /mnt. cd /mnt 3.Execute the following command to find directories occupying significant space. du –x ––max–depth=1 | sort –n Based on the returned results, assess whether directories or files can be deleted considering business requirements. 4.Execute the umount command to unmount the disk. For example: umount /mnt

Linux Instance Memory–Related Faults

Linux Instance Memory–Related Faults

Last Updated At: 2025–08–14 11:11:37

Symptom Description Faults caused by memory problems arise on Linux cloud server instances. Examples include slower internal service responses, inability to log into the server, and Out Of Memory (OOM) triggers in the system. Possible Causes Issues could be triggered by high memory utilization rates on the instance. Typically, when instance memory utilization remains above 90%, it can be judged as excessive memory usage. Troubleshooting Ideas 1.Follow the Handling Steps to determine if the issue is caused by high memory utilization. 2.Refer to Other Typical Cases of Memory Problems for further analysis of potential causes. Handling Steps 1.Check if the memory utilization is excessively high as described in Related Operations. o If memory utilization is high, proceed to the next step. o If memory utilization is normal, refer to Other Typical Cases of Memory Problems for additional cause identification. 2.After executing the top command inside the system and pressing M, check if any processes have high memory consumption under the ”RES” and ”SHR” columns. o If not, move to the next step. o If yes, operate on the corresponding process types for detailed instructions, see Analyzing Processes. 3.Execute the following command to check if shared memory usage is excessively high. cat /proc/meminfo | grep –i shmem 4.Execute the following command to check if unreclaimable slab memory usage is excessively high. cat /proc/meminfo | grep –i SUnreclaim 5.Execute the following command to check for existence of large pages in memory. cat /proc/meminfo | grep –iE ”HugePages_Total|Hugepagesize” o If HugePages_Total outputs 0, refer to Other Typical Cases of Memory Problems for further cause identification. o If HugePages_Total outputs non–zero values, it indicates that large pages in memory are configured. The size of memory large pages is HugePages_Total * Hugepagesize. You need to confirm whether hugepages were set– up by other malicious programs. If you determine that large pages in memory are no longer needed, you can disable– the

Error log reports: fork: Cannot allocate memory Phenomenon description The error message ”fork: Cannot allocate memory” appears in the logs. Possible reasons It may be due to exceeding process limits. When the total number of processes within the system reaches pid_max, creating a new– process will report a ”fork: Cannot allocate memory” error. Resolution approach 1.Refer to Troubleshooting Steps to check whether the instance’s memory usage rate is too high. 2.Verify if the total number of processes has exceeded the limit, and adjust the pid_max configuration accordingly. Troubleshooting steps 1.Refer to High Memory Usage Problem Handling to check if the instance has excessive memory usage. If the instance’s memory usage is normal, proceed to the next step. 2.Execute the following command to view the system’s pid_max value. sysctl –a | grep pid_max Based on the returned results, perform corresponding operations: □If the default pid_max value is 32768, proceed to the next step. □If the error message ”fork: Cannot allocate memory” is displayed, execute the following command to temporarily

increase pid_max. echo– 42768 > /proc/sys/kernel/pid_max. You can run the command again to view the system's pid_max value. 3.Execute the following command to view the total number of internal processes in the system. pstree –p | wc –l If the total number of processes reaches pid_max, the system will report a "fork: Cannot allocate memory" error when attempting to create a new process. Note You can execute the ps –efL command to identify programs that start many processes. 4.Modify the kernel.pid_max value in the /etc/sysctl.conf configuration file to 65535 to increase the number of processes. 5.Execute the following command to immediately apply the configuration changes. sysctl –p Out–of–memory triggered before exhausting instance memory Phenomenon description The Linux cloud server triggers OOM (Out of Memory) without fully utilizing its memory. Possible causes This could be caused by system available memory falling below the min_free_kbytes value. The min_free_kbytes value represents the minimum idle memory (in Kbytes) that the Linux system forcibly retains. If the system's available memory falls below the set min_free_kbytes value, the system defaults to starting oom–killer or forcing a reboot. The specific behavior is determined by the kernel parameter vm.panic_on_oom: If vm.panic_on_oom=0, the system will display an OOM warning and initiate oom–killer to terminate the process using the most memory. If vm.panic_on_oom=1, the system will automatically restart. Resolution approach 1.Follow the troubleshooting steps to investigate whether the instance's memory usage is excessively high or if the total number of processes is restricted. 2.Verify the setting of the min_free_kbytes value and modify it to the correct configuration. Troubleshooting steps 1.Refer to High Memory Usage Problem Handling to check if the instance has excessive memory usage. If the instance's memory usage is normal, proceed to the next step. 2.Refer to Error Log Reports: fork: Cannot Allocate Memory to verify if the number of processes has exceeded the limit. If the total number of processes has not exceeded the limit, proceed to the next step. 3.Log in to the cloud server and execute the following command to view the min_free_kbytes value. sysctl –a | grep min_free 4.Execute the following command to open the /etc/sysctl.conf configuration file using the VIM editor. vim /etc/sysctl.conf 5.Press i to enter edit mode and modify the vm.min_free_kbytes configuration item. If this configuration item does not exist, simply add it directly to the configuration file. Note It is recommended to set the vm.min_free_kbytes value to no more than 1% of the total memory. 6.Press Esc and then type: wq followed by Enter to save and exit the VIM editor. 7.Execute the following command to apply the configuration changes. sysctl —p

Network-related issues

Network-related issues

Last Updated At: 2025-08-14 14:00:06

Problem with multiple queue configurations for network interface cards Phenomenon description Misconfiguration of multiple queues for network interface cards on the cloud server. Possible reasons By default, the cloud server configures multiple queues for network interface cards, distributing NIC interrupts across different CPUs to improve network processing performance. Manual modifications might lead to errors in the multi-queue configuration for network interface cards. Resolution approach Follow the Troubleshooting Steps to correct the number of network card queues. Troubleshooting steps Assuming the default primary network interface card on the cloud server is eth0, and the number of network card queues is 2. 1.Execute the following command to view the current number of network card queues. ethtool -l eth0 Results indicating that the currently configured number of queues is less than the maximum number of network card queues supported by the NIC suggest an unreasonable setup requiring correction. Channel parameters for eth0: Pre-set maximums: RX: 0 TX: 0 Other: 0 Combined: 2 ### Maximum number of network card queues supported by the server Current hardware settings: RX: 0 TX: 0 Other: 0 Combined: 1 ### Currently configured number of network card queues 2.Execute the following command to set the current number of network card queues. ethtool -L eth0 combined 2 Adjust the queue count according to actual needs, setting the value equal to the maximum number of network card queues supported by the server. 3.Execute the following command to check the current configuration for the number of network card queues. ethtool -l eth0 If the maximum number of network card queues supported by the server matches the currently configured number of network card queues, the configuration was successful. CVM Network Access Packet Loss This document mainly introduces the main reasons that may cause CVM network access packet loss, as well as the corresponding troubleshooting and solutions. Possible Reasons The possible reasons for the CVM network access packet loss problem are as follows: Triggering the speed limit causes TCP packet loss. Triggering the speed limit causes UDP packet loss. Triggering soft interrupt causes packet loss UDP send buffer is full. UDP receive buffer is full. TCP full connection queue is full. TCP request overflow. The number of connections has reached the upper limit iptables policy sets relevant rules. Prerequisites You need to log in to the instance before locating and troubleshooting the problem. See Logging in to the Linux instance and Logging in to the Windows instance for details. Troubleshooting Triggering the speed limit causes TCP packet loss CVM instances are available in various specifications, and different specifications have different network performance. When the bandwidth or packet volume of an instance exceeds the standard corresponding to the instance specification, the speed limit on the platform side will be triggered, resulting in packet loss. The troubleshooting and processing steps are as follows: 1.Check the bandwidth and package size of the instance. For Linux instances, you can perform the sar -n DEV 2

command to check the bandwidth and packet volume. Among them, the metrics rxpck/s and txpck/s are the number of the sending and receiving packets, and the indicators rxkB/s and txkB/s are the sending and receiving bandwidth. 2.Use the obtained bandwidth and package data to compare instance specifications to check whether the instance specification performance bottleneck has been reached. –Yes, you need to upgrade the instance specifications or adjust the business volume. –No, if the performance bottleneck of the instance specification has not been reached, further location processing can be performed. Triggering the speed limit causes UDP packet loss See the step Triggering the speed limit causes TCP packet loss to determine whether the packet loss is caused by the performance bottleneck of the instance specification. □Yes, you need to upgrade the instance specifications or adjust the business volume. □If the performance bottleneck of the instance specification is not reached, it may be caused by the platform's additional frequency limit on DNS requests. When the overall bandwidth or packet volume of the instance reaches the performance bottleneck of the instance specifications, the DNS request speed limit may be triggered and UDP packet loss may occur. It can be further located and processed. Triggering soft interrupt causes packet loss When the operating system detects that the count value in the second column of /proc/net/softnet_stat is increasing, it will be judged as soft interrupt packet loss. When your instance triggers soft interrupt packet loss, you can troubleshoot and handle it by following the steps below: Check whether RPS is enabled: □If so, a small value of the kernel parameter net.core.netdev_max_backlog will cause packet loss and needs to be increased. □If not, check whether the CPU single–core soft interrupt is high, resulting in failure to send and receive data in time. If so, you can: –Enable RPS to make the soft interrupt distribution more balanced. –Check whether the business program causes uneven distribution of soft interrupts. UDP send buffer is full If your instance has packet loss due to insufficient UDP buffer, you can troubleshoot by following these steps: 3. 1.Use the ss –nump command to check whether the UDP send buffer is full. 2.If so, increase the kernel parameters net.core.wmem_max and net.core.wmem_default, and restart the UDP program for the changes to take effect. 3.If the packet loss problem still exists, you can use the ss –nump command to check whether the send buffer has not increased as expected. At this time, you need to check whether the business code has set SO_SNDBUF through setsockopt. If so, modify the code to increase SO_SNDBUF. UDP receive buffer is full If your instance has packet loss due to insufficient UDP buffer, you can handle it by following these steps: 4. 1.Use the ss –nump command to check whether the UDP receive buffer is full. 2.If so, increase the kernel parameters net.core.rmem_max and net.core.rmem_default, and restart the UDP program for the changes to take effect. 3.If the packet loss problem still exists, you can use the ss –nump command to check whether the receive buffer has not increased as expected. At this time, you need to check whether the business code has set SO_RCVBUF through setsockopt. If so, modify the code to increase SO_RCVBUF. TCP full connection queue is full The length of the TCP full connection queue is the smaller value of net.core.somaxconn and the backlog parameter passed in when the business process calls listen. If your instance has packet loss due to the full TCP connection queue, you can handle it by following these steps: 1.Increase the kernel parameter net.core.somaxconn. 2.Check whether the backlog parameter is passed into the business process. If so, increase it accordingly. TCP request overflow When TCP receives data, if

the socket is locked by the user, the data will be sent to the backlog queue. If this process fails, it will cause TCP request overflow and packet loss. Under normal conditions, assuming that the business program performance is normal, you can see the following methods to troubleshoot and solve the problem from the system level: Check whether the business program has set the buffer size by itself through setsockopt: □If it has been set and the value is not large enough, you can modify the business program to specify a larger value, or stop specifying the size through setsockopt. Description The value of setsockopt is limited by the kernel parameters net.core.rmem_max and net.core.wmem_max. When adjusting the business program, you can also adjust net.core.rmem_max and net.core.wmem_max simultaneously. After adjustment, restart the business program to make the configuration take effect. □If not set, you can increase the net.ipv4.tcp_mem, net.ipv4.tcp_rmem, and net.ipv4.tcp_wmem kernel parameters to adjust the TCP socket water level. The number of connections has reached the upper limit CVM instances are available in various specifications, and different specifications have different connection performance metrics. When the number of connections of an instance exceeds the standard corresponding to the instance specification, the platform's speed limit will be triggered, resulting in packet loss. The processing steps are as follows: Description The number of connections refers to the number of sessions of the CVM instance saved on the host, including TCP, UDP, and ICMP. This value is greater than the number of network connections obtained through the ss or netstat command on the CVM instance. Check the number of connections on your instance and compare it to Strength Specifications, check whether the instance specification performance bottleneck has been reached. □Yes, you need to upgrade the instance specifications or adjust the business volume. □No, if the performance bottleneck of the instance specification has not been reached, further location processing can be performed. iptables policy sets relevant rules In the absence of relevant iptables rules set on the CVM, the iptables policy–related rule settings may be causing packets destined for the CVM to be discarded. The processing steps are as follows: 4. 1.Perform the following command to check the iptables policy rules. iptables –L | grep policy The default iptables policy rule is ACCEPT. If the INPUT chain policy is not ACCEPT, all packets to the server will be lost. For example, if the following result is returned, it means that all packets entering the CVM will be dropped. Chain INPUT (policy DROP) Chain FORWARD (policy ACCEPT) Chain OUTPUT (policy ACCEPT) 2.Perform the following command and adjust the value after –P as needed. iptables –P INPUT ACCEPT After adjustment, you can perform the command in step 1 of iptables policy sets related ruless again to check the results as follows: Chain INPUT (policy ACCEPT) Chain FORWARD (policy ACCEPT) Chain OUTPUT (policy ACCEPT) Domain Name Cannot Be Resolved (CentOS 6.x system) Phenomenon Description After restarting the CVM with the operating system CentOS 6.x or performing the command service network restart, the CVM is unable to resolve domain names. At the same time, check the /etc–/resolv.conf configuration file, it was found that the DNS information was cleared. Possible Reasons In CentOS 6.x operating systems, the initscripts edition earlier than 9.03.49––1 have defects due to different grep editions. Ideas Upgrade the initscripts to the latest edition and regenerate the DNS information. Processing Procedures 1.Log in to the CVM. 2.Perform the following command to check the edition of initscripts and confirm whether the initscripts have defects due to an edition lower than

9.03.49--1. rpm –q initscripts Information resembles the following is returned: initscripts–9.03.40–2.e16.centos.x86_64 It can be seen that the initscripts edition initscripts–9.03.40--2 is lower than the problematic edition (initscripts–9.03.49--1), and there are risks of DNS being cleared. Perform the following command in sequence to upgrade initscripts to the latest edition and regenerate DNS information. yum makecache yum –y update initscripts service network restart 4.After the upgrade is complete, perform the following command to check the edition information of initscripts and confirm whether the upgrade is successful. rpm –q initscripts Information resembles the following is returned: initscripts–9.03.58–1.el6.centos.2.x86_64 It can be seen that the displayed edition is different from the previous edition and is higher than initscripts–9.03.49--1, and the upgrade operation is successful.

Common Issues

Storage

Storage

Last Updated At: 2025-08-14 14:26:29

System Disk What is the Default Space of the CVM System Disk? Currently, the default space of the system disk of a newly purchased CVM is 50 GB. Can the system disk of a cloud server be changed from a local disk to a cloud disk? • During application for a cloud server instance During application, you can directly select the disk type for the cloud server system disk. In Which Regions and Availability Zones Can the System Disk Be Resized to a Size Larger Than 50 GB? When the system disk is a CBS, all regions that support snapshots support adjusting the system disk to a size larger than 50 GB. When I Reinstall the System, Does the System Disk of a CVM Support Scaling Out (Capacity Expansion)? It can be divided into the following two scenarios. Please refer to them based on your actual situation: ●The system disk is a CBS: not supported. ●The system disk is a local hard disk: When you reinstall the system, there are two situations according to the size of the current system disk: -The default system disk space is 50 GB when you purchase the instance, and capacity scale-out is not supported. -This situation applies to instances purchased earlier: if the system disk space is less than or equal to 20 GB, it will be reinstalled to 20 GB by default; if the system disk space is greater than 20 GB, it will be reinstalled to 50 GB by default. Does the system disk support scaling out (expanding the capacity) and then scaling in (reducing the capacity) by reinstalling the system? The system disk does not support scale-in. When a small-capacity image of less than 50 GB is selected for creating or reinstalling a CVM, what is the size of the system disk? The selected small image does not affect the size of the system disk, and the minimum size is 50 GB. Does the cloud server system disk support partitioning? It is supported, but we do not recommend performing partitioning operations on the system disk. If you forcibly use third-party tools to partition the system disk, it may trigger unknown risks such as system crashes and data loss. Issues with using cloud disks How to View the Data Disk? 1.Log in to the CVM console. 2.Select CBS in the left sidebar to enter the CBS management page.

3.Click on the Properties column, select the data disk, click OK, and you can view all the data disks under the relevant region. After you reinstall the Windows system to the Linux system, how to read and write the original NTFS type data disk? Windows file systems usually use NTFS or FAT32 format, while Linux file system formats are usually EXT series. When the operating system is reinstalled from Windows to Linux, the operating system type changes, but the data disk remains in the original format. The reinstalled system may be unable to access the data disk file system. At this point, you need format conversion software to read the existing data. For specific operations, see Reading and Writing Original NTFS Type Data Disks After Reinstalling Windows as Linux. After Reinstalling the Linux System to the Windows System, How Do I Read the Original EXT-Type Data Disk? Windows file systems usually use NTFS or

FAT32 format, while Linux file system formats are usually EXT series. When the operating system is reinstalled from Linux to Windows, the operating system type changes, but the data disks remain in the original format. The reinstalled system may be unable to access the original data disk file system. At this point, you need format conversion software to read the existing data. For specific operations, see Reading Original EXT Type Data Disks After Reinstallation from Linux to Windows. What are the differences and similarities between various types of cloud disks? • Common Cloud Disk: This is the legacy cloud disk type provided by the cloud platform, suitable for low I/O load scenarios where data is rarely accessed. • High-Performance Cloud Disk: This is a hybrid storage type launched by the cloud platform, providing nearly solid-state storage level high-performance storage capabilities through cache mechanisms, while ensuring data reliability with a tri-replica distributed mechanism. High-performance cloud disks are suitable for small and medium applications with high data reliability requirements and average performance demands. • SSD Cloud Disk: Based entirely on NVMe SSD storage media, it adopts a tri-replica distributed mechanism to provide low-latency, high random IOPS, high throughput I/O capabilities, and data security at 99.9999999%, delivering high-performance storage. SSD cloud disks are suitable for scenarios with higher I/O performance requirements. Enhanced SSD Cloud Disk: This is a product type designed by the cloud platform based on a next-generation storage engine, utilizing full NVMe SSD storage media. It provides low-latency, high random IOPS, high throughput I/O capabilities, and data security at 99.9999999% through a tri-replica distributed mechanism. Enhanced SSD cloud disks are suitable for mid-sized databases, NoSQL, and other I/O-intensive scenarios with extremely high latency requirements. How can disk performance be tested? It is recommended to use FIO for stress testing and validation of cloud disks. How can the usage and remaining space of a cloud disk be checked? You can log in to the cloud server instance and view the usage and remaining space of the cloud disk internally within the instance. Alternatively, you can check the cloud disk usage situation through the Cloud Server Control Panel. Here are the steps: 1.Log in to the Cloud Server Control Panel and enter the 'Instances' list page. 2.Select the instance ID you want to check and enter the instance detail page. 3.On the instance detail page, select the Monitoring tab to view the cloud disk usage of that instance. Why was my separately created cloud disk released along with my instance? Cloud disks can be set to release automatically with the instance when detached. This can be achieved through the Cloud Disk Control Panel or by using APIs to alter the cloud disk attributes. Mounting and Unmounting Cloud Disk Issues What is a device name (mount point)? A device name (mount point) refers to the position of a cloud disk on the disk controller bus on a cloud server instance. The selected device name corresponds to the disk device number under the Linux operating system and aligns with the disk order in the disk manager under the Windows operating system. Can a single cloud disk be mounted onto multiple cloud server instances? Not currently supported. You can mount up to 20 cloud disks to a single cloud server, but currently, simultaneous sharing of a single cloud disk across multiple cloud servers is not supported. Data sharing can only be achieved by unmounting from Cloud Server A and then mounting to Cloud Server B. After applying for a cloud disk and mounting it to a cloud server instance, do I still need to execute the mounting and partitioning operation? After applying for a cloud disk, you need to mount it to a cloud

server in the same availability zone, and then perform initialization operations such as formatting, partitioning, and creating a file system before it can be used as a data disk. For specific operations, please refer to Mounting a Cloud Disk and Initializing a Cloud Disk. I applied for a data disk for a Linux instance, but I can't see it in the system. What should I do? If it's a separately applied–for data disk, you need to partition, format, and mount it before you can use and see the space. For specific operations, please refer to Mounting a Cloud Disk and Initializing a Cloud Disk. How many cloud disks can a single cloud server instance mount? As a data disk, a single instance can mount up to 20 cloud disks. Why can't I find the cloud server I want to mount the cloud disk to when mounting a cloud disk? Ensure that your cloud server instance has not been released and that both the cloud server instance and the cloud disk are in the same region and availability zone. Can a cloud disk and a cloud server instance in different availability zones be mounted together? No. You can only freely mount and unmount cloud disks among different cloud server instances within the same availability zone. Will the data on the cloud disk be lost when unmounting a cloud disk (data disk)? Data on the cloud disk does not change due to mounting or unmounting. To maintain data consistency, we suggest: • Under the Windows operating system, to ensure data integrity, it is recommended that you pause all file system read and write operations on the cloud disk, otherwise incomplete reads and writes will cause data loss. • Under the Linux operating system, you need to log in to the instance and run the umount command on the cloud disk. Once the command executes successfully, go to the control panel to unmount the cloud disk. Does the cloud disk support mounting and unmounting? • Cloud disks support mounting and unmounting. • System disks do not support mounting and unmounting. Does the cloud disk support batch mounting and unmounting? • Cloud disks support batch mounting and unmounting. • System disks do not support mounting and unmounting. Can the system disk be unmounted? No. Expansion and Reduction of Cloud Disk Issues How do I expand a cloud disk? When your cloud server is a cloud disk server, expansion can be performed. For operation guidelines, refer to Expanding a Cloud Disk. Can I compress the capacity of a cloud disk? The cloud platform does not support reducing the capacity of a cloud disk. If you need to reduce the capacity of a large cloud disk you've applied for, it is recommended that you create and mount a new cloud disk with appropriate capacity first, copy the required data from the old disk to the new disk, and then release the old disk. How do I expand the system disk? The cloud platform supports expanding system disk space through the cloud server control panel. For operation guidelines, please refer to Expanding a System Disk. Do all types of cloud disks support system disk expansion? SSD cloud disks, high–performance cloud disks, and common cloud disks support system disk expansion. Snapshot Usage Issues Are there any regional limitations for snapshots? At present, snapshot features are supported in all availability zones. Will taking a snapshot affect disk performance? Taking a snapshot occupies a small amount of I/O of the cloud disk. We recommend that you perform snapshot operations during relatively idle periods of your business. How long does it take to make a snapshot available after making a snapshot? The time taken to make a snapshot depends on factors such as the amount of data written to the cloud disk, underlying read and write conditions, and is difficult to predict. However, taking a snapshot does not affect your normal use of the disk. Do I need to shut down the instance when rolling back a snapshot? • For cloud

disks already mounted on a cloud server, shutting down the cloud server is necessary when rolling back. • For unmounted cloud disks, rollback operations can be executed directly. How is the capacity of the initial full snapshot of a cloud disk calculated? The first snapshot created for a cloud disk is a full snapshot, backing up all data on the cloud disk at a particular moment. The snapshot capacity equals the used capacity of the cloud disk. For example, if a cloud disk has a total capacity of 200GB and 122GB has been used, the size of the initial full snapshot would be 122GB. Can snapshots of cloud server instances be downloaded or exported locally? Snapshots cannot be downloaded or exported locally. Are there any differences or conflicts between manual snapshots and scheduled snapshots? There are no conflicts in usage, but timing conflicts might occur during creation. • When an automatic snapshot is being executed for a particular disk, the user needs to wait until the automatic snapshot completes before creating a custom snapshot (vice versa). • If the disk contains a large volume of data and the duration of a single snapshot exceeds the interval between two automatic snapshot time points, the next time point will skip the automatic snapshot. For example, if a user sets 9:00 AM, 10:00 AM, and 11:00 AM as automatic snapshot time points, and the execution of the 9:00 AM automatic snapshot takes 70 minutes (completing at 10:10 AM), then the 10:00 AM automatic snapshot will not be executed, and the next snapshot time point will be 11:00 AM. Do local disks support creating snapshots? No. It is recommended that you implement data redundancy at the application layer or create deployment sets for clusters to improve application availability. Will local snapshots be deleted along with the cloud disk upon releasing the cloud disk? No. If you need to delete related snapshots, please go to the control panel or use APIs to delete them. For details, please refer to Deleting Snapshots. Why is there inconsistency between the disk usage seen under the file system and the snapshot size? Cloud disk snapshots are block-level clone backups. Generally, snapshot capacities will be larger than the data volume statistics reported by the file system. The difference in capacity is caused by the following reasons: • Lower-level data blocks store metadata of the file system. • Deletion of data. Deleting data involves modifying written data blocks, and snapshots back up all modified data blocks. How do I retain snapshots to avoid deletion by the cloud platform? • Avoid arrears on your cloud platform account. If your account is in arrears, snapshots will enter an ”isolated” status. Snapshots in an ”isolated” status will be retained for 30 days. If the balance is not topped up to zero or greater during this period, all snapshots (except image snapshots) under the account will be deleted after the deadline. • Modify the retention time attribute of the scheduled snapshot strategy to long-term retention. When the automatic snapshots of a cloud disk reach the upper limit, the earliest created automatic snapshot will be automatically deleted. For detailed steps, please refer to Scheduled Snapshots. For information on snapshot quotas, see Usage Limitations. How do I delete snapshots to reduce backup usage costs? • For snapshots of cloud disks, they can be directly deleted through the control panel or via APIs. For specific operations, please refer to Deleting Snapshots. • For associated snapshots of custom images, you need to delete the corresponding custom image first before deleting the snapshot. After an instance expires or a cloud disk is released, will automatic snapshots be deleted? Automatic snapshots follow the retention time settings of scheduled snapshot strategies and will not be automatically deleted upon expiration of the instance or release of the cloud disk. If you need to modify

scheduled snapshot strategies, please refer to Scheduled Snapshots. How do I delete snapshots that have been used to create images or cloud disks? • Snapshots created for cloud disks can be deleted separately. After deleting a snapshot, you cannot operate businesses dependent on the original data state of the snapshot. • Snapshots created for custom images must have their corresponding images deleted beforehand in order to delete the snapshot. • Snapshots created for instances can be deleted separately. After deleting an image, you cannot operate businesses dependent on the original data state of the snapshot. If I use scheduled snapshots to create custom images or cloud disks, will executing snapshot policies fail? No. Can multiple automatic snapshot strategies be set for a single cloud disk? No. How can I avoid data loss caused by erroneous operations? In scenarios such as modifying critical system files, migrating instances from basic networks to private networks, routine data backups, accidental release/recovery of instances, preventing network attacks, changing operating systems, providing data support for production environments, and other operation–risk scenarios, you can Create Snapshots in advance to back up data. In case of erroneous operations, you can promptly Rollback Data from the snapshot to minimize risks. What's the difference between snapshots and images? Assuming an instance has no attached data disks and all data is written onto the system disk, creating only an image would not provide protection for this system disk. This is because images lack regular creation features; if the system disk data gets damaged, it can only trace back to the initial data at the time of image creation, which does not serve as adequate protection. The detailed differences are shown in the table below:

Name	Snapshot	Image
Nature	Data backup of cloud disk at a specific point in time	Template for software configuration of Cloud Server (operating system, pre–installed programs, etc.)
Applicable scenarios	□Regular backup of critical business data□Data backup before major operations□Multiple copies application of production data	□Backup systems that will not change in the short term□Bulk deployment of applications□System migration

How to migrate snapshot data from account A to account B? Snapshots cannot be migrated. If necessary, you can create an image from the snapshot and then share it under other accounts. Can data disk snapshots create custom images? No. The attribute of the cloud disk used to create a custom image must be a system disk.

Image

Image

Last Updated At: 2025-08-14 14:26:29

Private Image How to Handle the Failure of Creating a Custom Image for the Windows System? If the Windows system fails to create an image, check the following in order: 1.Make sure the following services and all our officially provided services starting with Win_Agent are running properly: 2.The execution of the custom image creation script is blocked by some antivirus tools or security dogs. To avoid creation failure, it is recommended to close these tools before creating a custom image. 3.The image creation tool was interrupted by a system pop-up window during execution. Log in to the CVM remotely to check and adjust the cloud server settings to avoid pop-ups. Can snapshots of data disks be used to create custom images? The disk attribute of the snapshot used to create a custom image must be a system disk; data disks cannot be used to create custom images. If you want to retain the data on the original instance's data disk when launching a new instance, you can first take a snapshot of the data disk and then use this snapshot to create a new cloud hard disk data disk when starting the new instance. For specific operations, please refer to Creating Cloud Hard Disks from Snapshots. How do you confirm that the data disk has been unmounted and can now be used to create a custom image? 1.Confirm that the line corresponding to the automatic mounting of the data disk partition in the /etc/fstab file has been deleted. 2.Use the mount command to view the mounting information of all devices, and ensure that the results do not contain any information about the corresponding data disk partitions.

Security

Security

Last Updated At: 2025-08-14 14:33:39

Password and Key What is the Difference Between SSH key Log-in and Password Log-in? SSH key is a way to remotely log in to a Linux server. Its principle is to use a key generator to create a pair of keys (public key and private key). Add the public key to the server, and then use the private key on the client to complete the authentication and login. This method pays more attention to data security, and is different from the manual input of the traditional password login method, and is more convenient. Currently, Linux instances support both password and SSH key log-in methods, while Windows instances only support password log-in at this time. Why can't I log in with a username and password after associating an SSH key with a Linux instance? After associating an SSH key with a cloud server, username and password login is disabled by default. Please log in to the cloud server using the SSH key. Can I Use SSH Keys to Log in and Passwords to Log in at the Same Time? Users log in to Linux instances using SSH keys, and password log-in is disabled by default to enhance security. Therefore, after logging in with a key, users will no longer be able to use a password to log in. How to Create SSH Keys and What to Do If You Lose Them? See Creating SSH Keys for how to create SSK keys. We provide two solutions to the key loss problem: □Create a new key through the SSH key console of the CVM and bind the new key to the original instance. i.Creating a SSH key. ii.After the key is created, log in to the CVM instance console. iii.Select the original instance to which the key is to be bound, click Operation > Password/Key > Load Key, and then you can use the new key to log in to the instance. □Reset your password through the CVM console and use the new password to log in to the instance. See Resetting Instance Password for more information. Why can't I download my key? Keys can only be downloaded once. If you've lost your key, we recommend recreating and downloading it for safekeeping. How can I check which key is being used for my cloud server instance? You can log in to the cloud server console, go to the details page of your cloud server instance, where you will find the key information associated with that instance. Security group related Why Is There a Default Deny Rule in the Security Group? Security group rules are applied sequentially from top to bottom. After the preceding allow rules are matched, other rules are denied by default. If all ports are open, the last deny rule will not take effect. For security reasons, we provide this default setting. What impact does selecting the incorrect security group have on instances bind to it? How Can This Be Resolved? Issue risks ●Remote connections (SSH) to Linux instances, and remote desktop logs in to Windows instances, may fail. ●Remote ping to the public IP and private network IP of CVM instances under this security group may fail. ●HTTP access to web services exposed by instances under this security group may fail. ●Instances under this security group may fail to failures access Internet services. Solution ●If the above problems occur, you can reset the security group rules in the security group management of the console. For example, only bind the default full connection security group. ●See

Security Group Introduction for more information about setting security group rules. What are the Direction and Policy of Security Groups? The direction of security group policies is divided into outbound and inbound. Outbound direction refers to filtering the outbound traffic of cloud servers, while inbound direction refers to filtering the inbound traffic of cloud servers. Security group policies are divided into Allow and Deny traffic. What is the Order of Effectiveness for Security Group Policies? From top to bottom. The matching order of traffic through security group policies is from top to bottom. Once a match is successful, the policy takes effect. Why can IPs That Are Not Allowed by the Security Group Still Access CVMs? There may be the following reasons: ●The CVM may bind to multiple security groups, and the specific IP is allowed in other security groups. ●The specific IP belongs to approved Cloud Platform public services. Does using a security group mean that iptables cannot be used? No. Security groups and iptables can be used simultaneously. Your traffic will be filtered twice, and the traffic flow is as follows: ●Outbound direction: processes on the instance > iptables > security group. ●Inbound direction: security group > iptables > processes on the instance. Can a Cloned Security Group Be Named the Same as the Security Group in the Target Region? No. The name must be different from the existing security group name in the target region. Does the Security Group Support Cross-user Cloning? Not support. Will cloning security groups across projects and regions also copy the cloud servers managed by the security group? No, when cloning security groups across regions, only the original security group's inbound and outbound rules are cloned, and cloud servers need to be associated separately. What is a security group? A security group acts as a virtual firewall with stateful packet filtering capabilities, used to set network access controls for instances such as cloud servers, load balancers, and cloud databases, controlling inbound and outbound traffic at the instance level, serving as an important means of network security isolation. Each cloud server instance belongs to at least one security group, and specifying a security group is mandatory when creating an instance. Cloud server instances within the same security group communicate freely over the network, whereas instances in different security groups are isolated by default, requiring explicit permission for inter-group communication. For more details, see Security Group Overview. Why choose a security group when creating a cloud server instance? Before creating a cloud server instance, you must select a security group to define the security domain for your application environment and authorize security group rules for proper network security segmentation. What if no security group was created prior to launching a cloud server instance? If you did not create a security group before initiating a cloud server instance, you can opt to create a new security group during the setup process. Creating a new security group offers the following rules, adjust according to your actual needs to allow specific IP/port combinations: • ICMP: Allow ICMP protocol, enabling public internet ping to the server. • TCP:80: Permit port 80, allowing HTTP access to web services. • TCP:22: Enable port 22, permitting SSH remote connection to Linux cloud servers. • TCP:443: Open port 443, allowing HTTPS access to web services. • TCP:3389: Grant access to port 3389, enabling RDP remote connection to Windows cloud servers. • Intranet Access: Allow intranet access, facilitating internal network connectivity between different cloud resources (IPv4).

Maintenance and Monitoring Related Common Operations and Commands in Linux

Last Updated At: 2025-08-14 14:41:01

What is the Load Average on a Linux Server? Load Average is a measure of how busy a computer system is, indicating the average system load over recent past periods, typically represented as a three-number moving average (one, five, and fifteen minutes ago). How do I view the load on a Linux server? To view the load on a Linux server, you can execute commands such as w, top, uptime, or access the contents of the /proc/loadavg file. For instructions on installing the procinfo tool in a Linux environment, please refer to relevant software installation documentation. What should I do if the server load is high? High server load (Load Average) might suggest various issues, including insufficient CPU resources, I/O read/write bottlenecks, lack of memory resources, or CPU-intensive computations. When observing high load, consider referencing the 15-minute average as an indicator. Utilize commands like vmstat, iostat, top to diagnose the underlying cause of excessive load and optimize resource consumption. How do I check the server's memory usage? To monitor the server's memory usage, you can execute commands such as free, top (after pressing Shift+M to sort by memory), vmstat, or inspect the /proc/meminfo file. How do I check the amount of memory used by a single process? To determine the memory consumed by a single process, you can execute commands like top -p PID, pmap -x PID, ps aux|grep PID, or examine the /proc/process_id/(replaceprocess_id)/status with the PID of the target process, for instance, /proc/7159/status. How do I view active services and ports? To list active services along with their associated ports, you can execute commands such as netstat -tunlp, netstat -antup, or lsof -i:PORT. How do I view information about server processes? To display information about server processes, you can execute commands such as ps auxww|grep PID, ps -ef, lsof -p PID, or top -p PID. How do I stop a process? To halt a process, you can execute kill -9 PID (where PID signifies the process identifier) or killall followed by the name of the program (for example, killall cron). Stopping Zombie Processes If aiming to terminate zombie processes, it's necessary to kill the parent process. This can be achieved by executing kill -9 ppid (where ppid denotes the parent process ID, which can be identified by running ps -o ppid PID, for example, ps -o ppid 32535). How do I find zombie processes? To locate zombie processes, you can execute the top command to view the aggregate count of zombie processes or run ps -ef | grep defunct | grep -v grep to pinpoint specific zombie processes. Why can't I start the server port? Troubleshooting inability to initiate a server port involves examining both the operating system and the application side. On Linux Operating Systems Ports below 1024 can exclusively be activated by the root user. Hence, before commencing the service port, you must execute sudo su — to obtain root-level privileges. Application-Specific Troubleshooting For issues originating from applications, it's advisable to investigate failure causes through application startup logs. Common culprits include port conflicts (such as occupying port

36000 utilized by the cloud platform server system), misconfigurations, among others. Below is the table displaying the common Linux server performance monitoring commands:

Command Name	Description
top	A process monitoring command, used to monitor the overall system performance.Displays system load, processes, CPU, memory, paging, etc. Commonly use Shift+M and Shift+P to sort processes by memory and CPU usage respectively.
vmstat	A system monitoring command, emphasizing virtual memory monitoring, also capable of monitoring CPU, processes, memory paging, and IO status information.For example, vmstat 3 10, outputs results every 3 seconds, repeated 10 times.
iostat	Tool for outputting CPU status and IO status, providing detailed information about system IO.For example, iostat –dxmt 10, outputs detailed IO information every 10 seconds in MB format.
top	Text
top	Text
top	Text

NTP Service Related

Last Updated At: 2025-08-14 15:03:00

After configuring the NTP service, how do you adjust the synchronization interval for NTP? After setting up the NTP service, you can reset the NTP synchronization interval by restarting the ntpd service. If you need to manually set the ntpd synchronization interval, follow these steps: 1.Execute the following command to modify the NTP configuration file. vi /etc/ntp.conf 2.Press i to enter edit mode and perform the following configurations: i. If there is server time1.yun.com iburst, add a # at the beginning of the line to comment it out. ii. Add the following configuration, where minpoll 4 indicates a minimum of 16 (which is 2^4), and maxpoll 5 indicates a maximum of 32 (which is 2^5). server time1.yun.com minpoll 4 maxpoll 5 After making changes, type :wq to save the modifications and exit.

Why was the ntp.conf content reverted after creating a cloud server using a custom image? This is due to the system's Cloud-Init initialization. Before creating a custom image, please remove the NTP-related configuration from /etc/cloud/cloud.cfg. What specific impacts would changing the internal DNS have?

System–related

System–related

Last Updated At: 2025–08–14 15:03:00

During shutdown or restart operations for cloud services, there is a very slim chance of failure. In case of failure, you can troubleshoot and handle the situation according to the following guidelines. Possible reasons for failure during shutdown/restart: 1.Check the CPU and memory usage of your cloud server. When CPU utilization is excessively high or memory is depleted, it may lead to failures in shutting down or restarting from the control panel. 2.For Linux operating systems, check whether the ACPI management program is installed by running the command ps –ef | grep –w ”acpid” | grep –v ”grep” to see if there is any process present. If not, install the acpid module. 3.For Windows operating systems, check if the prolonged presence of WindowsUpdate leads to failed shutdowns because when doing certain patch operations, Windows performs some processing during system shutdown, which may result in lengthy update times causing failure to shut down or restart. 4.When Windows is initially applied for, due to the distribution of images using Sysprep, the initialization process takes slightly longer. Before completion of initialization, Windows will ignore shutdown/restart operations leading to failed attempts at shutting down or restarting. 5.If certain software has been installed on the operating system or if it has been infected with malware or viruses, the system itself might be compromised, potentially resulting in failures to shut down or restart. Force Shutdown/Restart Functionality: The cloud platform offers forced shutdown/restart functionality. In situations where multiple attempts to shut down or restart your cloud server have failed, you can use this feature. This operation forcefully shuts down or restarts your cloud server and could lead to data loss or file system damage on the cloud server. Select ”force shutdown” in the cloud server control panel’s shutdown operation Select ”force restart” in the cloud server control panel’s reboot operation.

Network and DNS

Network and DNS

Last Updated At: 2025-08-14 15:03:00

Domain Name Cannot Be Resolved (CentOS 6.x system) Phenomenon Description After restarting the CVM with the operating system CentOS 6.x or performing the command service network restart, the CVM is unable to resolve domain names. At the same time, check the /etc-/resolv.conf configuration file, it was found that the DNS information was cleared. Possible Reasons In CentOS 6.x operating systems, the initscripts edition earlier than 9.03.49--1 have defects due to different grep editions. Ideas Upgrade the initscripts to the latest edition and regenerate the DNS information. Processing Procedures 1.Log in to the CVM. 2.Perform the following command to check the edition of initscripts and confirm whether the initscripts have defects due to an edition lower than 9.03.49--1. rpm -q initscripts Information resembles the following is returned: initscripts-9.03.40-2.e16.centos.x86_64 It can be seen that the initscripts edition initscripts-9.03.40--2 is lower than the problematic edition (initscripts-9.03.49--1), and there are risks of DNS being cleared. 3.Perform the following command in sequence to upgrade initscripts to the latest edition and regenerate DNS information. cat /dev/null>/etc-/resolv.conf yum makecache yum -y update initscripts service network restart 4.After the upgrade is complete, perform the following command to check the edition information of initscripts and confirm whether the upgrade is successful. rpm -q initscripts Information resembles the following is returned: initscripts-9.03.58-1.el6.centos.2.x86_64 It can be seen that the displayed edition is different from the previous edition and is higher than initscripts-9.03.49--1, and the upgrade operation is successful. Unable to ping the instance IP address Failure Symptoms The local host's failure to ping the instance could be due to the following problems: the target server settings are incorrect, the domain name failed to resolve properly, Link Fault. On the premise that the local network is functioning properly (you can ping other websites without issue), you can troubleshoot according to the following steps: 1.Check whether the instance is configured with a public IP The instance must have a public IP address to be accessible to other computers on the Internet. If the instance does not have a public IP, the private IP address cannot directly ping the instance. 2.Check security group settings A security group is a virtual firewall that controls inbound and outbound traffic for associated instances. The rules of the security group can specify protocols, ports, policies, etc. Because ping uses the ICMP protocol, confirm whether the security group associated with instance allows ICMP. The security group used by the instance and detailed inbound and outbound rules can be viewed on the Security Group tab of the instance details page. 3.Check system settings Check Linux kernel parameters and firewall settings Whether the Linux system allows ping is determined by both the kernel and firewall settings. If either of them blocks the system, it will cause a Request timeout for ping packets. Check the value of kernel parameter icmp_echo_ignore_all icmp_echo_ignore_all indicates whether the system ignores all ICMP Echo requests, where 1 represents prohibition and 0 represents permission. Use the following

command to view the system’s icmp_echo_ignore_all setting: cat /proc/sys/net/ipv4/icmp_echo_ignore_all You can modify it using the echo command:

Solution for Unable to Create Network Namespace Issues

Solution for Unable to Create Network Namespace Issues

Last Updated At: 2025-08-14 15:17:41

Problem Description When executing a command to create a new network namespace (NetworkNamespace), the command gets stuck and fails to continue. dmesg message: ”unregister_netdevice: waiting for lot to become free. Usage count = 1”. Cause of Problem This is a kernel bug. Currently, the following kernel versions contain this bug: – Ubuntu 16.04 x86_64 kernel version 4.4.0-91-generic; – Ubuntu 16.04 x86_32 kernel version 4.4.0-92-generic. Resolution Upgrade the kernel version to 4.4.0-98-generic, which has fixed this bug. Directions Check the current kernel version. uname -r Check if version 4.4.0-98-generic is available for upgrade. sudo apt-get update sudo apt-cache search linux-image-4.4.0-98-generic Displaying the following information indicates that the source contains this version, allowing for an upgrade: linux-image-4.4.0-98-generic – Linux kernel image for version 4.4.0 on 64-bit x86 SMP Install the new kernel version and the corresponding header package. sudo apt-get install linux-image-4.4.0-98-generic linux-headers-4.4.0-98-generic Restart the system. sudo reboot Enter the system and check the kernel version. uname -r Showing the following result indicates successful version update: 4.4.0-98-generic

General Problems

General Problems

Last Updated At: 2025-08-14 14:05:58

Investigating High CPU Utilization (LINUX) High CPU utilization can lead to slower service responses and difficulties logging into servers. You can utilize Cloud Monitoring to establish CPU utilization threshold alerts, which notify you promptly when CPU usage exceeds the defined thresholds. The general steps for investigating high CPU utilization involve identifying the specific processes consuming CPU resources and analyzing those with high CPU occupancy. If identified as abnormal processes, they may be virus or malware related; you can either manually terminate these processes or use security software to eliminate them. For business processes, analyze whether increased traffic volume is causing the spike and determine if optimization is possible. In case of platform component processes, submit a ticket for us to conduct further investigation and handling. Below we explain how to pinpoint high CPU utilization processes under the Linux system. Introduction to Diagnostic Tools: top Command top: A commonly used monitoring tool in Linux systems for acquiring real-time CPU usage at the process level. The upper half displays overall resource usage of CPU and memory: First line: Current system time, number of logged-in users, and system load. Second line: Total number of system processes, running processes, sleeping, hibernated, and zombie process counts. Third line: Current CPU usage. Fourth line: Current memory usage. Fifth line: Current swap space usage. The lower half shows resource occupation by process dimension. PID: Process ID. USER: Owner of the process. PR: Priority of the process NI: NICE value, where smaller values indicate higher priority. VIRT: Size of virtual memory used, measured in KB. RES: Current size of physical memory used, measured in KB. SHR: Size of shared memory used, measured in KB. S: State of the process. %CPU: Percentage of CPU time used by the process during the update interval. %MEM: Percentage of memory used by the process during the update interval. TIME+: CPU time consumed by the process, accurate to 0.01 seconds. COMMAND: Name of the process. Issue Identification and Resolution Using Tools to Identify Processes with High CPU Utilization We've previously introduced the top tool; now let's discuss how to leverage it to pinpoint processes with high CPU utilization. Log in to the instance via SSH. 1.Enter the top command to view system load. 2.Press uppercase P to sort processes by CPU utilization in descending order; through sorting, easily identify processes heavily consuming CPU resources for further analysis. 3.Analyze processes with high CPU consumption. (1)For business processes, consider optimizing program efficiency or upgrading instance resources. (2)For abnormal processes, suspect potential infection; terminate the process yourself, use antivirus software for elimination, or after data backup, reinstall the system. (3)For cloud platform component processes occupying over 20% CPU, submit a ticket for us to conduct further investigation and handling. Common cloud platform components include: sap00x: Security component process Barad_agent: Monitoring component process secu-tcs-agent: Security component process Investigating High CPU Utilization

(WINDOWS) High CPU utilization can lead to slowed service responses and difficulties logging into servers. You can use Cloud Monitoring to set up CPU utilization threshold alerts, notifying you instantly when CPU usage surpasses specified levels. Investigating high CPU utilization involves identifying specific processes consuming CPU resources and analyzing those with high CPU occupancy. If categorized as abnormal processes, likely due to viruses or malware, you can manually terminate these processes or use security software for eradication; for business processes, analyze whether increased traffic volumes are responsible and determine if optimization is feasible; if platform component processes, submit a ticket for further investigation and handling. Below we detail how to pinpoint high CPU utilization problems under Windows systems. Introduction to Diagnostic Tools Task Manager: A built-in application and process management tool in Windows, displaying information about computer performance and running software, including names of running processes, CPU loads, memory usage, I/O situations, logged-in users, and Windows services. Accessible via shortcut keys Ctrl+Shift+Esc, right-clicking Start Menu and selecting Task Manager, or typing taskmgr in Run dialog box. Processes: List of all processes running on the system. Performance: Overall statistics regarding system performance such as aggregate CPU usage and utilized memory. Users: All users currently having sessions on the system. Details: An enhanced version of the Processes tab showing detailed information about processes such as their PID, status, CPU/memory usage, etc. Services: All services in the system (including inactive ones). Issue Identification and Resolution High CPU utilization might be caused by hardware factors, system processes, business processes, or Trojan viruses among other elements. Below, we describe how to pinpoint the specific processes consuming CPU resources and how to analyze and manage these processes. Log into the Windows server. 1.Open Task Manager using Ctrl+Shift+Esc or right-clicking the Start Menu and choosing Task Manager, switch to the Details tab, click on CPU to sort processes by CPU utilization in descending order. 2.Analyze processes with significant CPU usage. These could be system, business, or anomalous processes; below we exemplify how to handle each scenario: (1)System processes. Upon discovering system processes consuming substantial CPU resources, carefully scrutinize the process name since malware often mimics system process names to deceive users. Examples include svch0st.exe, explore.exe, iexplorer.exe; ensure careful verification. Also, check the location of executable files associated with these processes; system processes typically reside in c:\windows\system32 and carry complete signatures and descriptions. Right-click the corresponding process in Task Manager and select ‘Open File Location’ to view the exact position of the executable file. If the process location isn’t under c:\windows\system32, the server may be infected; manually or through security tools, eradicate any threats. Common system processes include: SystemIdleProcess (system idle process, displaying percentage of CPU idle time), system (memory management process), explorer (desktop and file management), iexplore (Microsoft browser), csrss (Microsoft client/server runtime subsystem), svchost (system process for executing DLLs), Taskmgr (Task Manager), lsass (local security authority service), etc. (2)Abnormal processes. If unusually named processes consume large amounts of CPU resources, they might be malware processes. Use search engines to confirm identification, e.g., xmr64.exe (mining virus). Once confirmed, employ security tools for elimination. (3)Business processes. Should your business processes (such as iis, httpd, php, java,

etc.) be identified as consuming CPU resources, further analysis is advised. Consider whether heavy current business volumes justify high loads; if so, contemplate upgrading server configurations. Otherwise, assess potential optimization opportunities for business programs.

Region and Availability Zone (AZ)

Region and Availability Zone (AZ)

Last Updated At: 2025-08-14 14:05:58

How to Check the Region List? You can check it in the following ways: ●Check the documentation for Region and Availability Zone. ●Querying through API: -Querying region list. -Querying availablity zone list. What regions and availability zones are available for cloud servers? How to choose? For available regions and availability zones of cloud servers, you may refer to Regions and Availability Zones. For choosing regions and availability zones, you may refer to How to Choose Regions and Availability Zones. Can the Purchased CVM be Changed to a Different Region? The region of the purchased CVM cannot be changed. The availability zone of an already launched instance cannot be changed. If you need to change the region and availability zone, see the following solutions: First refund the instance, then reapply for a new instance. Create a custom image of the original instance first, then use the custom image to create an instance in the new availability zone, start the instance, and update the configuration of the new instance. i.Create a custom image for the current instance. See Creating a Custom Image for more information. ii.If the network environment of the current instance is VPC and the current private IP address needs to be retained after migration, you can first delete the subnet in the current availability zone and then create a subnet in the new availability zone with the same IP address range as the original subnet. Note: If the deleted subnet contains available instances, move all instances in the current subnet to the new subnet and delete them. iii.Create an instance in the new availability zone using the custom image you just created.

Instance–related

Instance–related

Last Updated At: 2025–08–14 14:10:41

Unable to log in remotely due to security group settings This document introduces troubleshooting methods and solutions for remote connection issues caused by security group settings on cloud servers. Step 1: Connection Test 1.On your local computer, press and hold both the ”Windows” key and the ”R” key simultaneously; type ”cmd” in the pop–up window and hit Enter to open Command Prompt.