ALIBABA CLOUD

Alibaba Cloud

Virtual Private Cloud
Product Introduction

Document Version: 20201016

Alibaba Cloud

Legal disclaimer

Alibaba Cloud reminds you to carefully read and fully understand the terms and conditions of this legal disclaimer before you read or use this document. If you have read or used this document, it shall be deemed as your total acceptance of this legal disclaimer.

1. You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloud-authorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.

2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company or individual in any form or by any means without the prior written consent of Alibaba Cloud.

3. The content of this document may be changed because of product version upgrade, adjustment, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and an updated version of this document will be released through Alibaba Cloud-authorized channels from time to time. You should pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.

4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides this document based on the "status quo", "being defective", and "existing functions" of its products and services. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not take legal responsibility for any errors or lost profits incurred by any organization, company, or individual arising from download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, take responsibility for any indirect, consequential, punitive, contingent, special, or punitive damages, including lost profits arising from the use or trust in this document (even if Alibaba Cloud has been notified of the possibility of such a loss).

5. By law, all the contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of this document shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates.

6. Please directly contact Alibaba Cloud for any errors of this document.

Document conventions

Style	Description	Example
Danger	A danger notice indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	Danger: Resetting will result in the loss of user configuration data.
Warning	A warning notice indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	Warning: Restarting will cause business interruption. About 10 minutes are required to restart an instance.
Notice	A caution notice indicates warning information, supplementary instructions, and other content that the user must understand.	Notice: If the weight is set to 0, the server no longer receives new requests.
Note	A note indicates supplemental instructions, best practices, tips, and other content.	Note: You can use Ctrl + A to select all files.
>	Closing angle brackets are used to indicate a multi-level menu cascade.	Click **Settings**> **Network**> **Set network type**.
Bold	Bold formatting is used for buttons , menus, page names, and other UI elements.	Click **OK**.
`Courier font`	Courier font is used for commands	Run the `cd /d C:/window` command to enter the Windows system folder.
Italic	Italic formatting is used for parameters and variables.	`bae log list --instanceid` *Instance_ID*
[] or [a|b]	This format is used for an optional value, where only one item can be selected.	`ipconfig [-all|-t]`
{} or {a|b}	This format is used for a required value, where only one item can be selected.	`switch {active|stand}`

Table of Contents

1.What is a VPC?

A virtual private cloud (VPC) is a private network dedicated for your use. You have full control over your VPC, which you can define and customize by specifying the Classless Inter-domain Routing (CIDR) block, configuring route tables, and creating gateways. You can launch Apsara Stack resources such as Elastic Compute Service (ECS) instances, ApsaraDB for RDS (RDS) instances, and Server Load Balancer (SLB) instances in your VPC.

Furthermore, you can connect your VPC to other VPCs or on-premises networks to create a custom network environment. In this way, you can smoothly migrate applications and extend on-premises data centers to the cloud.

Components

Each VPC consists of one VRouter, at least one private CIDR block, and one or more VSwitches.

- Private CIDR block

  When you create a VPC or a VSwitch, you must specify its private IP address range in the form of a CIDR block.

  You can use the standard private CIDR blocks listed in the following table and their subsets as CIDR blocks for your VPCs. For more information, see Set up network connections.

CIDR block	Number of available private IP addresses (excluding those reserved by the system)
192.168.0.0/16	65,532
172.16.0.0/12	1,048,572
10.0.0.0/8	16,777,212

- VRouter

  A VRouter is a hub that connects all VSwitches in a VPC and serves as a gateway between the VPC and other networks. After a VPC is created, a VRouter is automatically created for the VPC. Each VRouter is associated with a route table.

  For more information, see Overview.

- VSwitch

  A VSwitch is a basic network component that connects different cloud resources in a VPC. After you create a VPC, you can create VSwitches to partition your VPC into multiple subnets. VSwitches within a VPC can communicate with each other over the private network. You can deploy your applications in VSwitches that belong to different zones to improve service availability.

  For more information, see VSwitches.

2.VPC connections

Alibaba Cloud provides a wide range of solutions to help you connect your VPC to other VPCs, the Internet, or on-premises data centers.

Connect a VPC network to the Internet

The following table lists the services that you can use to connect a Virtual Private Cloud (VPC) network to the Internet.

Service	Feature	Benefit
ECS public IP address	When you create an Elastic Compute Service (ECS) instance in a VPC network, you can allow the system to automatically assign a public IP address to the ECS instance. Then, the ECS instance can use the public IP address to communicate with the Internet. You cannot unbind a public IP address from an ECS instance when the ECS instance is running. However, you can convert the public IP address to an elastic IP address (EIP). For more information, see Convert an automatically assigned public IP address to an EIP for a VPC-connected ECS instance.	You can purchase data transfer plans for an ECS instance that is assigned public IP addresses. You can also purchase EIP bandwidth plans for an ECS instance after you convert the public IP address of the ECS instance to an EIP. For more information, see What is EIP bandwidth plan and What is a data transfer plan.
Elastic IP Address	EIPs can be associated with or disassociated from ECS instances anytime. ECS instances can use EIPs to communicate with the Internet based on Source Network Address Translation (SNAT) and Destination Network Address Translation (DNAT).	EIPs can be associated with or disassociated from ECS instances anytime. You can use EIP bandwidth plans and data transfer plans to reduce the costs of data transfer over the Internet.
NAT Gateway	You can create SNAT and DNAT entries on a NAT gateway to enable one or more ECS instances in a VPC network to communicate with the Internet. Note Unlike Server Load Balancer (SLB), NAT gateways are incapable of balancing the load of ECS instances.	A NAT gateway can provide Internet access for more than one ECS instance while an EIP can serve only one ECS instance.

Service	Feature	Benefit
Server Load Balancer	Server Load Balancer (SLB) is a port-based service that provides Layer-4 and Layer-7 load balancing. ECS instances that are connected to SLB can be accessed over the Internet. Note SLB does not support SNAT. ECS instances deployed in VPC networks cannot access the Internet through SLB.	SLB supports DNAT. Each port on an SLB instance can be mapped to one or more ECS instances. SLB distributes traffic to ECS instances to balance the load of the ECS instances. This improves the availability of application systems and eliminates single points of failure. After you associate an EIP with an SLB instance, you can purchase EIP bandwidth plans and data transfer plans to reduce costs.

Connect VPC networks

The following table lists the services that you can use to connect two VPC networks.

Service	Feature	Benefit
Cloud Enterprise Network (CEN)	Connects VPC networks in different regions or under different accounts. For more information, see 教程概览.	• Global connection. • Low latency and high speed. • Nearby access and shortest-path connection. • Connection redundancy and disaster recovery. • Systematic management.
VPN Gateway	Connects two VPC networks through an IPsec-VPN connection for encrypted data transmission. For more information, see Establish a connection between two VPCs.	• Enhanced security. • High availability. • Cost-effectiveness. • Easy configuration.

Connect a VPC network to an on-premises data center

The following table lists the services that you can use to connect a VPC network to an on-premises data center.

Service	Feature	Benefit
Express Connect	Connects a VPC network to an on-premises data center. For more information, see What is physical connection.	• Low-latency data transmission through the backbone network. • Higher security and reliability based on leased lines.

Service	Feature	Benefit
VPN Gateway	• Connects an on-premises data center to a VPC network through an IPsec-VPN connection. • Connects a client to a VPC network through an SSL-VPN connection.	• Enhanced security. • High availability. • Cost-effectiveness. • Easy configuration.
CEN	• Connects a VPC network to an on-premises data center. You can build an interconnected network by adding the virtual border router (VBR) that is associated with an on-premises data center to a CEN instance. • Connects multiple VPC networks to an on-premises data center. You can build an interconnected network by adding multiple network instances (VPC/VBR) to a CEN instance.	• Global connection. • Low latency and high speed. • Nearby access and shortest-path connection. • Connection redundancy and disaster recovery. • Systematic management.
Smart Access Gateway (SAG)	• Connects on-premises networks, such as data centers and branches, to Alibaba Cloud to build a hybrid cloud. • Connects on-premises networks.	• Supports automatic configurations and zero touch provisioning (ZTP), and automatically adapts to network topology changes. • Connects to nearby access points in a metropolitan area network. On-premises networks can be connected to Alibaba Cloud through primary and secondary connections/devices. • Builds a secure hybrid cloud. Data transmission among VPC networks and over the Internet is encrypted.

3.Architecture

Based on the tunneling technique, VPCs isolate virtual networks. Each VPC has a unique tunnel ID, and each tunnel ID corresponds to only one VPC.

Background information

With the development of cloud computing, a variety of network virtualization techniques have been developed to meet the increasing demands for virtual networks with higher scalability, security, reliability, privacy, and connectivity.

Earlier solutions combined the virtual network with the physical network to form a flat network, for example, the large layer-2 network. However, with the increase of virtual network scale, problems such as ARP spoofing, broadcast storms, and host scanning are becoming more serious. To resolve these problems, various network isolation techniques are developed to completely isolate the physical network from the virtual network. One of these techniques can isolate users with a VLAN. However, a VLAN only supports up to 4,096 nodes, which are insufficient for the large number of users in the public cloud.

Principles

Based on the tunneling technique, VPCs isolate virtual networks. Each VPC has a unique tunnel ID, and each tunnel ID corresponds to only one VPC. A tunnel encapsulation carrying a unique tunnel ID is added to each data packet transmitted over the physical network between ECS instances in a VPC. In different VPCs, ECS instances with different tunnel IDs are located on two different routing planes. Therefore, these ECS instances cannot communicate with each other.

Based on the tunneling and Software Defined Network (SDN) techniques, Alibaba Cloud has developed VPCs that are integrated with gateways and VSwitches.

Logical architecture

As shown in the following figure, a VPC consists of a gateway, a controller, and one or more VSwitches. The VSwitches and gateway form a key data path. By using a protocol developed by Alibaba Cloud, the controller distributes the forwarding table to the gateway and VSwitches to provide a key configuration path. In the overall architecture, the configuration path and data path are separated from each other. The VSwitches are distributed nodes, the gateway and controller are deployed in clusters, and all links are equipped with disaster recovery. These features improve the availability of the VPC.

4.Benefits

This topic describes the benefits of using VPCs.

High security

Each VPC has a unique tunnel ID, and each tunnel ID corresponds to a virtual network. Different VPCs are isolated by tunnel IDs:

- Similar to traditional networks, VPCs can also be divided into subnets. ECS instances in the same subnet use the same VSwitch to communicate with each other, while ECS instances in different subnets use VRouters to communicate with each other.
- VPCs are completely isolated from each other and can only be interconnected by mapping an EIP or a NAT IP address.
- ECS IP packets are encapsulated by using the tunneling technique. Therefore, information about the data link layer (layer-2 MAC address) of ECS does not go to the physical network. As a result, the layer-2 network between different ECS instances or between different VPCs is isolated.
- ECS instances in a VPC use security groups as firewalls to control traffic going to and from ECS instances. This is layer-3 isolation.

High flexibility

You can use security groups or whitelists to flexibly control traffic going to and from the cloud resources in a VPC.

Ease of use

You can quickly create and manage VPCs in the VPC console. After a VPC is created, the system automatically creates a VRouter and a route table for the VPC.

High scalability

You can create multiple subnets in a VPC to deploy different services. Additionally, you can connect a VPC to other VPCs or on-premises data centers to expand your network.

5.Scenarios

This topic describes the scenarios in which VPCs are used to guarantee a high level of data security and service availability.

Host applications that provide external services

You can host applications that provide external services in a VPC and control access to these applications from the Internet by creating security group rules and access control whitelists. You can also isolate Internet-based mutual access between the application server and the database. For example, you can deploy the web server in a subnet that can access the Internet and deploy the application database in a subnet that cannot access the Internet.

Host applications that require access to the Internet

You can host applications that require access to the Internet in a subnet of a VPC and route traffic through network address translation (NAT). After you configure SNAT rules, instances in the subnet can access the Internet without exposing their private IP addresses, which can be changed to public IP addresses any time to avoid external attacks.

Implement disaster tolerance across zones

You can create one or multiple subnets in a VPC by creating VSwitches. VSwitches in a VPC can communicate with each other. You can deploy resources on VSwitches in different zones for disaster tolerance.

Isolate business systems

VPCs are logically isolated from each other. Therefore, you can create multiple VPCs to isolate multiple business systems, for example, isolate the production environment from the test environment. You can also create a peering connection between two VPCs if they need to communicate with each other.

Build a hybrid cloud

You can create a dedicated connection to connect your VPC to an on-premises data center to expand your local network. By doing so, you can seamlessly migrate your local applications to the cloud without changing the method of access to these applications.

Control big bandwidth fluctuations caused by multiple applications

If your applications generate big bandwidth fluctuations, you can configure DNAT forwarding rules through the NAT Gateway. Then, you can add EIPs to Internet Shared Bandwidth so that these EIPs can share the bandwidth. This can reduce bandwidth fluctuations and save your cost.

6.Terms

This topic describes the terms about VPCs.

Term	Description
Virtual Private Cloud (VPC)	A VPC is a private network established in Alibaba Cloud. VPCs are logically isolated from each other. You can create and manage cloud resources in your VPC, such as ECS, SLB, and RDS.
VSwitch	A VSwitch is a basic network device that connect different cloud resources in a VPC. When you create a cloud resource in a VPC, you must specify the VSwitch to which the cloud resource is connected.
VRouter	A VRouter is a hub that connects all VSwitches in a VPC and serves as a gateway that connects the VPC to other networks. A VRouter also forwards network traffic according to the route entries in its route table.
Route table	A route table is a list of route entries in a VRouter.
Route entry	Each item in a route table is a route entry. A route entry specifies the next hop address for the network traffic directed to a destination CIDR block. Route entries are divided into system route entries and custom route entries.

7.Limits

Before you use a virtual private cloud (VPC), note the following limits.

Limits on VPCs and VSwitches

Item	Default limit	Quota increase
Number of VPCs that can be created in each region	10	Go to the Quota Management page and request a quota increase. For more information, see Quota management.
Number of VSwitches that can be created in each VPC	24	
Available CIDR blocks for each VPC	192.168.0.0/16, 172.16.0.0/12, 10.0.0.0/8, and their subnets	To use a public CIDR block as the VPC CIDR block,submit a ticket.
Number of secondary IPv4 CIDR blocks that can be created in each VPC	1	
Number of user CIDR blocks that can be created in each VPC	3	

Item	Default limit	Quota increase
Number of private IP addresses that can be used by cloud resources in each VPC	60,000 **Note** • If an Elastic Compute Service (ECS) instance has only one private IP address, the ECS instance must use the IP address for communication. • If an ECS instance is associated with one or more elastic network interfaces (ENIs), or the ENI of the ECS instance is assigned multiple IP addresses, the number of IP addresses that the ECS instance can use equals the total number of IP addresses assigned to the ENIs.	N/A
Number of tags that can be attached to each VPC	20	
Number of tags that can be attached to each VSwitch	20	

Limits on VRouters and route tables

Item	Default limit	Quota increase
Number of VRouters that can be created in each VPC	1	N/A

Item	Default limit	Quota increase
Number of custom route tables that can be created in each VPC	9 Note Custom route tables are supported in all regions except China (Beijing), China (Shenzhen), and China (Hangzhou).	Go to the Quota Management page and request a quota increase. For more information, see Quota management.
Number of custom route entries that can be created in each route table	48	
VPCs that do not support custom route tables	VPCs that contain ECS instances of the following instance families: ecs.c1, ecs.c2, ecs.c4, ecs.ce4, ecs.cm4, ecs.d1, ecs.e3, ecs.e4, ecs.ga1, ecs.gn4, ecs.gn5, ecs.i1, ecs.m1, ecs.m2, ecs.mn4, ecs.n1, ecs.n2, ecs.n4, ecs.s1, ecs.s2, ecs.s3, ecs.se1, ecs.sn1, ecs.sn2, ecs.t1, and ecs.xn4. For more information, see Overview of VPC advanced features.	N/A
Number of tags that can be attached to each route table	20	

Limits on shared VPC

Item	Default limit	Quota increase
Number of participants supported by each VPC	20	
Number of participants supported by each VSwitch in a VPC	20	
Number of VSwitches that can be shared with each participant	10	

Item	Default limit	Quota increase
Number of IP addresses that each VPC can use	Shared by the resource owner and participants	N/A
Types of VSwitches that can be shared	Non-default VSwitches	
Regions that support shared VPCs	• Singapore (Singapore) • China (Zhangjiakou-Beijing Winter Olympics) • China (Hangzhou) • China (Shanghai)	
Cloud resources that can be created in a shared VSwitch	• ECS instances • SLB instances • RDS instances	
Limits on security groups in a shared VPC	• A participant cannot create resources with security groups that belong to other participants or the resource owner, including the default security group • The resource owner cannot create resources with security groups that belong to participants	

Limits on flow logs

Item	Default limit	Quota increase
Number of flow logs that can be created in each region	10	N/A

Item	Default limit	Quota increase
VPCs that do not support flow logs	VPCs that contain instances of the following instance families: ecs.c1, ecs.c2, ecs.c4, ecs.ce4, ecs.cm4, ecs.d1, ecs.e3, ecs.e4, ecs.ga1, ecs.gn4, ecs.gn5, ecs.i1, ecs.m1, ecs.m2, ecs.mn4, ecs.n1, ecs.n2, ecs.n4, ecs.s1, ecs.s2, ecs.s3, ecs.se1, ecs.sn1, ecs.sn2, ecs.t1, and ecs.xn4.	Upgrade or release an Elastic Compute Service (ECS) instance that does not support advanced network features. • For more information, see Upgrade configurations of subscription instances and Change the instance type of a pay-as-you-go instance. • For more information, see Release an instance. Note If the VPC to which a VSwitch or ENI belongs contains one of the specified instance families and the flow logs feature is enabled, you must upgrade or release the instance for flow logs to work as expected. For more information, see Overview of VPC advanced features.
VSwitches that do not support flow logs	VPCs to which VSwitches belong contain instances of the following instance families: ecs.c1, ecs.c2, ecs.c4, ecs.ce4, ecs.cm4, ecs.d1, ecs.e3, ecs.e4, ecs.ga1, ecs.gn4, ecs.gn5, ecs.i1, ecs.m1, ecs.m2, ecs.mn4, ecs.n1, ecs.n2, ecs.n4, ecs.s1, ecs.s2, ecs.s3, ecs.se1, ecs.sn1, ecs.sn2, ecs.t1, and ecs.xn4.	
ENIs that do not support flow logs	VPCs to which ENIs belong contain instances of the following instance families: ecs.c1, ecs.c2, ecs.c4, ecs.ce4, ecs.cm4, ecs.d1, ecs.e3, ecs.e4, ecs.ga1, ecs.gn4, ecs.gn5, ecs.i1, ecs.m1, ecs.m2, ecs.mn4, ecs.n1, ecs.n2, ecs.n4, ecs.s1, ecs.s2, ecs.s3, ecs.se1, ecs.sn1, ecs.sn2, ecs.t1, and ecs.xn4.	

Limits on network access control lists (ACLs)

Item	Default limit	Quota increase
Number of network ACLs that can be created in each VPC	200	N/A
Number of network ACLs that can be associated with each VSwitch	1	

Item	Default limit	Quota increase
Number of rules that can be added to a network ACL	• Inbound rules: 20 • Outbound rules: 20	Go to the Quota Management page and request a quota increase. For more information, see Quota management.
VPCs that do not support network ACLs	VPCs that contain instances of the following instance families: ecs.c1, ecs.c2, ecs.c4, ecs.ce4, ecs.cm4, ecs.d1, ecs.e3, ecs.e4, ecs.ga1, ecs.gn4, ecs.gn5, ecs.i1, ecs.m1, ecs.m2, ecs.mn4, ecs.n1, ecs.n2, ecs.n4, ecs.s1, ecs.s2, ecs.s3, ecs.se1, ecs.sn1, ecs.sn2, ecs.t1, and ecs.xn4. For more information, see Overview of VPC advanced features.	Upgrade or release an Elastic Compute Service (ECS) instance that does not support advanced network features. • For more information, see Upgrade configurations of subscription instances and Change the instance type of a pay-as-you-go instance. • For more information, see Release an instance. Note If the VPC contains one of the specified ECS instance families and the network ACL feature is enabled, you must upgrade or release the ECS instance for the network ACL to work as expected.

Limits on high-availability virtual IP addresses (HAVIPs)

Item	Default limit	Quota increase
Network types that support HAVIPs	VPCs	
Number of HAVIPs that can be created under each account	5	
Number of HAVIPs that can be created in each VPC	5	
Number of HAVIPs that can be associated with each ECS instance	5	
Number of ECS instances that can be associated with each HAVIP	2	

Item	Default limit	N/A Quota increase
Number of route entries destined for an HAVIP in each VPC	5	
Whether HAVIPs support broadcast or multicast communication	N/A Note HAVIPs support only unicast. To implement high availability through third-party software such keepalived, you must modify the configuration file to change the communication method to unicast.	