

Enterprise Switch

Service Overview

Issue 01
Date 2023-11-17



Copyright © Huawei Cloud Computing Technologies Co., Ltd. 2024. All rights reserved.

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Cloud Computing Technologies Co., Ltd.

Trademarks and Permissions



HUAWEI and other Huawei trademarks are the property of Huawei Technologies Co., Ltd.

All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

The purchased products, services and features are stipulated by the contract made between Huawei Cloud and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Huawei Cloud Computing Technologies Co., Ltd.

Address: Huawei Cloud Data Center Jiaoxinggong Road
Qianzhong Avenue
Gui'an New District
Gui Zhou 550029
People's Republic of China

Website: <https://www.huaweicloud.com/intl/en-us/>

Contents

1 What Is an Enterprise Switch?	1
2 How Enterprise Switches Work	2
3 Why Using Enterprise Switches	6
4 Notes and Constraints	9
5 Permissions Management	12
6 Region and AZ	14
7 Working with Other Services	16
8 Billing	17

1 What Is an Enterprise Switch?

Enterprise switches enable Layer 2 networking for VPCs, helping you to connect cloud and on-premises networks that are highly reliable, in a large scale, and of high performance.


Currently, enterprise switches only support Layer 2 connection gateways (L2CGs). An L2CG is a virtual tunnel gateway that can work with Direct Connect or VPN to establish network communications between cloud and on-premises networks at Layer 2. The gateway allows you to migrate workloads in data centers or private clouds to the cloud without changing subnets and IP addresses.

VPN or Direct Connect only allows cloud and on-premises networks to communicate at Layer 3 and the CIDR blocks of the networks that are used for communication cannot overlap.

If the cloud and on-premises networks overlap and need to communicate with each other, you can use an enterprise switch to enable communication between them at Layer 2.

An enterprise switch is a tunnel gateway of a VPC and corresponds to the tunnel gateway of your data center. It can work together with Direct Connect or VPN to enable communications between a VPC and your data center at Layer 2. **Figure 1-1** shows the networking diagram. You need to connect a VPC subnet to the enterprise switch and specify the enterprise switch to establish a connection with the tunnel gateway of your on-premises data center so that the VPC subnet can communicate with the data center subnet at Layer 2.

Figure 1-1 Layer 2 networking



2 How Enterprise Switches Work

Figure 2-1 illustrates how an enterprise switch works. Table 2-1 describes the working principles in more detail.

Figure 2-1 Networking

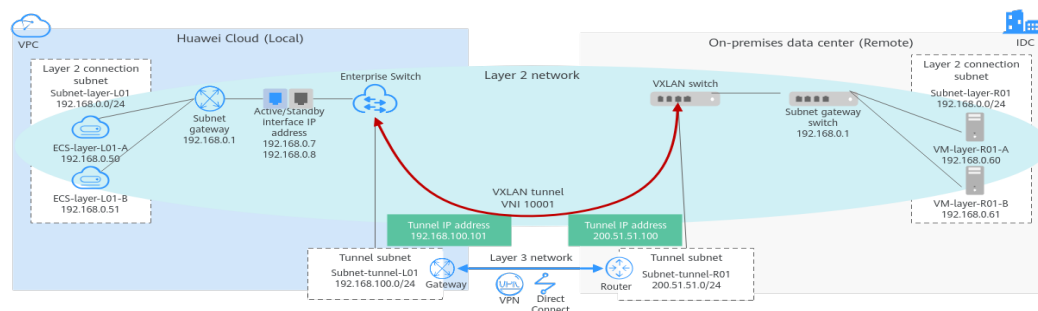


Table 2-1 Working principles

No.	Action	Description
1	Enable the local and remote tunnel subnets to communicate at Layer 3.	<ul style="list-style-type: none"> Plan resources on and off the cloud. For details, see Table 2-2. Use Direct Connect or VPN to enable the local (Subnet-tunnel-L01) and remote (Subnet-tunnel-R01) tunnel subnets to communicate at Layer 3.
2	Create an enterprise switch and specify a tunnel subnet .	Create an enterprise switch, set the local tunnel subnet to Subnet-tunnel-L01 , and the local tunnel IP address to 192.168.100.101 .

No.	Action	Description
3	Create a Layer 2 connection .	Create a Layer 2 connection to enable the local Layer 2 connection subnet (Subnet-layer-L01) and the remote VXLAN switch to communicate at Layer 2. Configure the following parameters: <ul style="list-style-type: none"> Active and standby interface IP addresses: They can be automatically assigned or manually specified. Remote tunnel IP address (200.51.51.100) and tunnel VNI (10001)
4	Configure a tunnel gateway in the on-premises data center.	Configure a tunnel gateway on the remote VXLAN switch to establish a VXLAN tunnel for the remote Layer 2 connection subnet (Subnet-layer-R01).

Table 2-2 Resource details

Resource	Cloud (Local)		On Premises (Remote)	
Layer 2 connection subnet	VPC subnet	Subnet-layer-L01: 192.168.0.0/24	On-premises subnet	Subnet-layer-R01: 192.168.0.0/24
	ECS	<ul style="list-style-type: none"> ECS-layer-L01-A: 192.168.0.50 ECS-layer-L01-B: 192.168.0.51 	On-premises server	<ul style="list-style-type: none"> VM-layer-R01-A: 192.168.0.60 VM-layer-R01-B: 192.168.0.61
	Active/Standby interface IP address	<ul style="list-style-type: none"> Active interface IP address: 192.168.0.7 Standby interface IP address: 192.168.0.8 	-	-
Tunnel subnet	VPC subnet	Subnet-tunnel-L01: 192.168.100.0/24	On-premises subnet	Subnet-tunnel-R01: 200.51.51.0/24
	Tunnel IP address	192.168.100.101	Tunnel IP address	200.51.51.100
Tunnel VNI	10001			

Layer 2 Connection Subnets

A local Layer 2 connection subnet is on the cloud and a remote one is in an on-premises data center. They are used to communicate at Layer 2.

- Local Layer 2 connection subnet: a VPC subnet, for example, Subnet-layer-L01
- Remote Layer 2 connection subnet: an on-premises subnet, for example, Subnet-layer-R01

Constraints

- The local and remote Layer 2 connection subnets can overlap, but the IP addresses of the servers that need to communicate in the local and remote subnets must be different. Otherwise, the communication fails.
- A VPC subnet that has been used a Layer 2 connection cannot be used by any other Layer 2 connections or enterprise switches.

Tunnel Subnets

Local and remote tunnel subnets communicate with each other at Layer 3 over Direct Connect or VPN. Enterprise switches allow communications between cloud and on-premises networks at Layer 2 based on the Layer 3 network between tunnel subnets.

- Local tunnel subnet: a VPC subnet, for example, Subnet-tunnel-L01
- Remote tunnel subnet: an on-premises subnet, for example, Subnet-tunnel-R01

Constraints

- Ensure that the local and remote tunnel subnets can communicate at Layer 3 over VPN or Direct Connect before you use an enterprise switch to allow communication at Layer 2.
- The switch in an on-premises data center must support VXLAN because the enterprise switch needs to establish a VXLAN tunnel to the data center at Layer 2.
- The local tunnel subnet must have three IP addresses reserved for the enterprise switch.

Layer 2 Connections

After an enterprise switch is created, you need to create a Layer 2 connection to enable the local Layer 2 connection subnet and the remote VXLAN switch to communicate at Layer 2.

Constraints

- Each Layer 2 connection connects a local and a remote Layer 2 connection subnet. Each enterprise switch supports a maximum of six Layer 2 connections.
- The Layer 2 connections of an enterprise switch can share a tunnel IP address, but their tunnel VNIs must be unique. A tunnel VNI is the identifier of a tunnel.
- If a Layer 2 connection connects a local Layer 2 connection subnet to an enterprise switch, the local Layer 2 connection subnet must have two IP

addresses reserved as active and standby interface IP addresses. The two IP addresses cannot be used by any local resources and must be different from the IP addresses in the remote Layer 2 connection subnet.

Active and Standby Interface IP Addresses

If a Layer 2 connection connects a local Layer 2 connection subnet to an enterprise switch, the local Layer 2 connection subnet must have two IP addresses reserved as active and standby interface IP addresses.

Tunnel IP Addresses

If an enterprise switch establishes a VXLAN tunnel with an on-premises data center at Layer 2, each end of the VXLAN tunnel requires a tunnel IP address, that is the local and remote tunnel IP addresses. The two IP addresses must be different.

- Local tunnel IP address: in the local tunnel subnet. In this example, the local tunnel subnet is Subnet-tunnel-L01, and the tunnel IP address is 192.168.100.101.
- Remote tunnel IP address: in the remote tunnel subnet. In this example, the remote tunnel subnet is Subnet-tunnel-R01, and the tunnel IP address is 200.51.51.100.

Tunnel VNIs

Tunnel VNIs are used to uniquely identify the VXLAN tunnels between an on-premises data center and an enterprise switch.

For the same VXLAN tunnel, the on-premises data center and the cloud must use the same tunnel VNI.

3 Why Using Enterprise Switches

VPN or Direct Connect allows communications between on-premises data centers and VPCs at Layer 3. However, this may require network reconstruction, long cloud migration period, and service interruptions. For details, see [Constraints on Communication at Layer 3](#).

Enterprise switches allow communications between on-premises data centers and VPCs at Layer 2, helping you dynamically and smoothly migrate workloads to the cloud. For details, see [Advantages on Communication at Layer 2](#).

Constraints on Communication at Layer 3

[Figure 3-1](#) shows the Layer 3 network between on-premises data centers and VPCs using VPN or Direct Connect. [Table 3-1](#) describes the pain points.

Figure 3-1 Layer 3 networking diagram

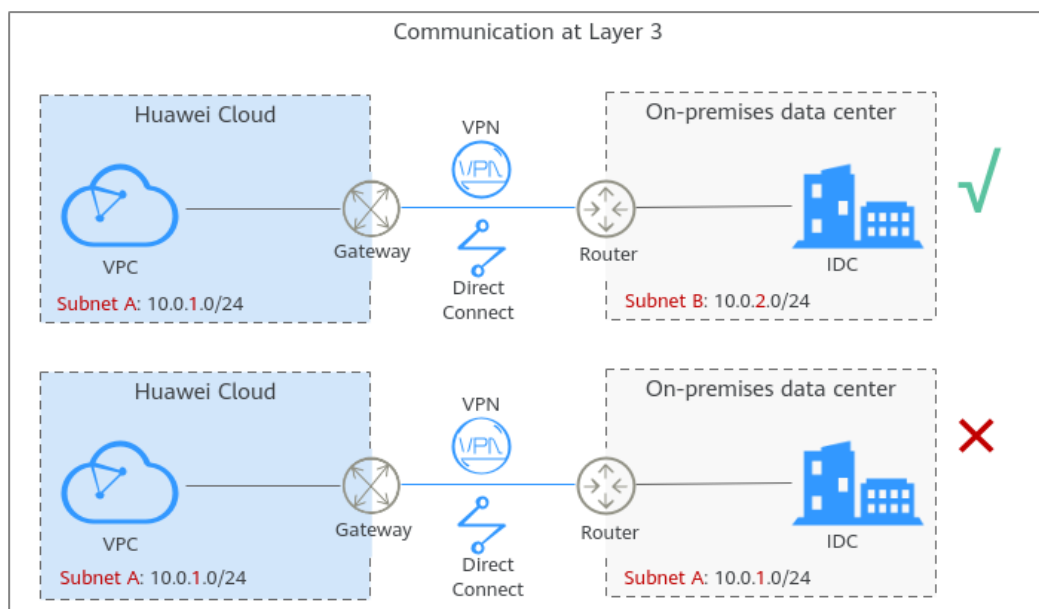


Table 3-1 Layer 3 networking description

Description	VPN or Direct Connect allows the communication between on-premises data centers and VPCs at Layer 3 through routes.
Pain Points	<ul style="list-style-type: none"> The CIDR blocks of the on-premises data center and the VPC that are used for communication cannot overlap. On-premises workloads communicate with each other using IP addresses instead of domain names. If the CIDR blocks of the on-premises data center and the VPC that are used for communication overlap, the on-premises network needs to be reconstructed before the cloud migration, which prolongs the cloud migration period, interrupts businesses, and increases O&M costs. Workloads in a subnet have to be migrated together, and cloud and on-premises workloads in the same subnet cannot communicate with each other. Dozens of different workloads are deployed on each subnet of the on-premises data center. If workloads are migrated by subnet, business continuity cannot be ensured.

Advantages on Communication at Layer 2

To handle the pain points of cloud migration at Layer 3, you can use enterprise switches to allow the communication between on-premises data centers and VPCs at Layer 2. For details about the advantages of enterprise switches, see [Table 3-2](#).

Figure 3-2 Layer 2 networking diagram

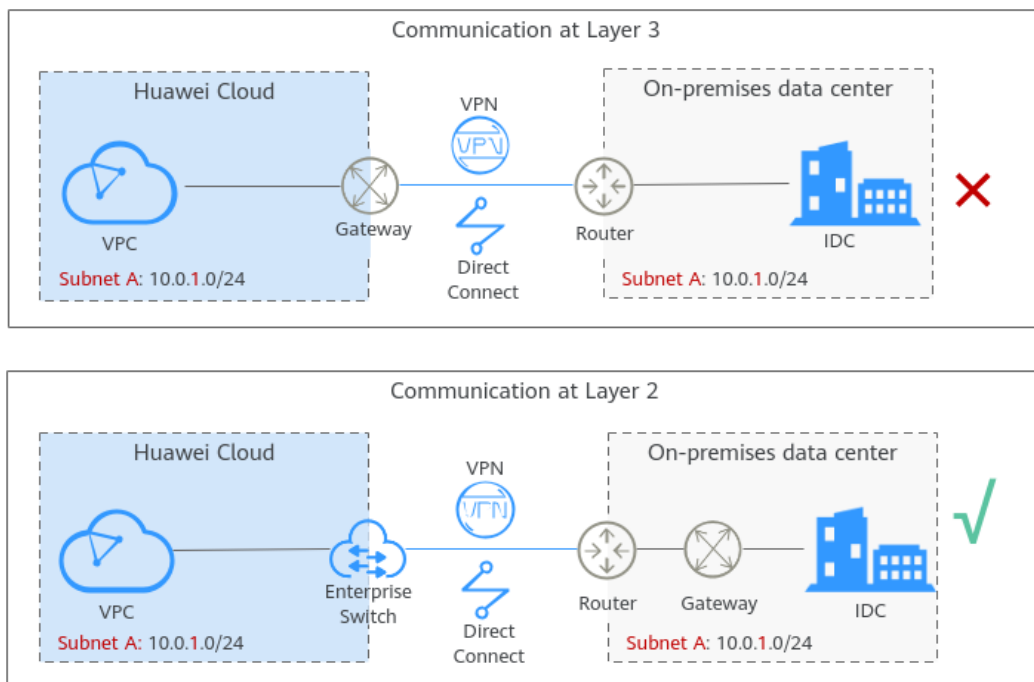


Table 3-2 Layer 2 networking description

Description	Enterprise switches establish a Layer 2 network between on-premises data centers and VPCs based on the Layer 3 network established by VPN or Direct Connect.
Advantages	<ul style="list-style-type: none">• The CIDR blocks of the on-premises data center and the VPC that are used for communication can overlap. An enterprise switch allows the network of the on-premises data center to remain unchanged even if the data center and the VPC have overlapping CIDR blocks.• Workloads can be migrated to the cloud on a server basis, and cloud and on-premises workloads in the same subnet can communicate with each other. Workloads can be seamlessly migrated to the cloud to prevent any loss caused by cloud migration.

4 Notes and Constraints

Quota Limits

Table 4-1 Quota limits

Resource	Default Quota	Adjustable
Maximum number of enterprise switches per account	5	Yes
Maximum number of VPCs that can be associated with an enterprise switch	1	No
Maximum number of Layer 2 connections supported by each VPC subnet	1	No
Maximum number of Layer 2 connections supported by each basic enterprise switch	1	No
Maximum number of Layer 2 connections supported by each standard enterprise switch	3	No
Maximum number of Layer 2 connections supported by each enhanced enterprise switch	6	No

Constraints

- Enterprise switches cannot forward unknown unicast, broadcast, and multicast (except VRRP) packets from your data center to the cloud or IPv6 packets.
- On-premises servers cannot use advanced network functions on the cloud, such as VPC Peering, Route Table, ELB, and NAT Gateway.

- If you want to use a Direct Connect connection together with an enterprise switch, [submit a service ticket](#) to check whether your connection can interconnect with an enterprise switch. If your connection does not support this, contact customer service.
- If you want to use a VPN connection together with an enterprise switch, [submit a service ticket](#) to check whether your connection supports VXLAN interconnection with an enterprise switch. If your connection does not support this, contact customer service.
- Only classic VPNs can be used together with enterprise switches.
- If cloud and on-premises networks communicate with each other at Layer 2, the cloud subnet gateway address must be the same as the on-premises subnet gateway address. Otherwise, the on-premises subnet gateway address may conflict with the IP address of a cloud server, causing communication exceptions.
- Each enterprise switch allows up to 10,000 IP addresses to communicate at Layer 2, including a maximum of 1,000 on-premises IP addresses.
- To use an enterprise switch to connect cloud and on-premises networks at Layer 2, a customer is responsible for constructing the VXLAN network of their data center equipment room, including preparing VXLAN switches, connecting physical networks, and connecting to Direct Connect or VPN.
- An enterprise switch supports MAC proxy forwarding. With ARP proxy, cloud and on-premises servers can communicate without knowing the actual MAC address of each other. The source MAC address of the packets received by the cloud server is the MAC address of the main interface of the Layer 2 connection, and the source MAC address of the packets received by the on-premises server is the MAC address of the tunnel interface. If your services need to know the actual server MAC addresses or have MAC address-based security policies, do not use an enterprise switch.
- Generally, a server determines the destination MAC address of a reply packet through ARP. However, some hosts or hardware devices (such as F5 load balancers) are configured to use the source MAC address of a request packet as the destination MAC address of its reply packet. If an enterprise switch is used for cloud and on-premises communications at Layer 3 in this case, a network disconnection may occur.

For example, an enterprise switch allows cloud and on-premises networks to communicate on 192.168.3.0/24. If the cloud server 192.168.2.2/24 accesses the on-premises server 192.168.3.3/24, the request packet on the cloud is sent to the on-premises server through the VPC route and then the enterprise switch. The reply packet from the on-premises server is sent back to the cloud through the route and Direct Connect or VPN. If the on-premises server is configured to use the source MAC address of a request packet as the destination MAC address of its reply packet, the destination MAC address of the reply packet is not the MAC address of gateway 192.168.3.0/24, but the source MAC address of the request packet, that is, the MAC address of the enterprise switch. In this case, the destination MAC address of the reply packet is incorrect and the network is disconnected.

- If an enterprise switch uses the VXLAN protocol, the VXLAN protocol has a header of 50 bytes and the packet length increases. Your on-premises network devices that allow VXLAN packets should support jumbo frames (Ethernet frames whose MTU is greater than 1500 bytes). Otherwise, such packets cannot be transmitted.

 **NOTE**

Vendors such as Huawei, allow jumbo frames by default. However, vendors such as Cisco, do not allow jumbo frames by default.

- If you use an enterprise switch to connect your on-premises data center to the cloud, the switch of your data center must support the VXLAN function. The following lists some switches that support the VXLAN function.
 - Huawei switches: Huawei CE58, CE68, CE78, and CE88 series switches, such as CE6870, CE6875, CE6881, CE6863, and CE12800 switches
 - Switches of other vendors: Cisco Nexus 9300, Ruijie RG-S6250, and H3C S6520 series switches

5 Permissions Management

If you need to assign different permissions to employees in your enterprise to control their access to your cloud resources, you can use Identity and Access Management (IAM) for fine-grained permissions management. IAM provides functions such as identity authentication, permissions management, and access control.

With IAM, you can create IAM users and assign permissions to the users to control their access to specific resources.

If your HUAWEI ID does not need individual IAM users for permissions management, skip this section.

IAM can be used free of charge. You pay only for the resources in your account. For more information about IAM, see [IAM Service Overview](#).

Enterprise Switch Permissions

By default, new IAM users do not have any permissions assigned. You need to add them to one or more groups and attach roles to these groups so that these users can inherit permissions from the groups and perform specified operations on cloud services.

Enterprise Switch is a project-level service deployed and accessed in specific physical regions. You need to select a project such as **ap-southeast-2** for which the permissions will be granted. If you select **All projects**, the permissions will be granted for all the projects. You need to switch to the authorized region before accessing Enterprise Switch.

Enterprise Switch uses the same system permissions as VPC. [Table 5-1](#) lists all the system-defined roles and policies supported by VPC. This VPC role is dependent on other roles. When assigning VPC roles to users, you need to also assign dependent roles for the VPC permissions to take effect.

Table 5-1 System-defined permissions for VPC

Policy Name	Description	Policy Type	Dependencies
VPC FullAccess	Full permissions for VPC	System-defined policy	To use the VPC flow log function, users must also have the LTS ReadOnlyAccess permission.
VPC ReadOnlyAccess	Read-only permissions on VPC.	System-defined policy	None
VPC Administrator	Most permissions on VPC, excluding creating, modifying, deleting, and viewing security groups and security group rules. To be granted this permission, users must also have the Tenant Guest permission.	System-defined role	Tenant Guest policy, which must be attached in the same project as VPC Administrator .

Helpful Links

- [IAM Service Overview](#)
- [Creating a User and Granting Permissions to Access the Enterprise Switch](#)

6 Region and AZ


Concept

A region and availability zone (AZ) identify the location of a data center. You can create resources in a specific region and AZ.

- Regions are divided based on geographical location and network latency. Public services, such as Elastic Cloud Server (ECS), Elastic Volume Service (EVS), Object Storage Service (OBS), Virtual Private Cloud (VPC), Elastic IP (EIP), and Image Management Service (IMS), are shared within the same region. Regions are classified into universal regions and dedicated regions. A universal region provides universal cloud services for common tenants. A dedicated region provides specific services for specific tenants.
- An AZ contains one or more physical data centers. Each AZ has independent cooling, fire extinguishing, moisture-proof, and electricity facilities. Within an AZ, computing, network, storage, and other resources are logically divided into multiple clusters. AZs within a region are interconnected using high-speed optical fibers, to support cross-AZ high-availability systems.

Figure 6-1 shows the relationship between regions and AZs.

Figure 6-1 Regions and AZs



Huawei Cloud provides services in many regions around the world. You can select a region and an AZ based on requirements. For more information, see [Huawei Cloud Global Regions](#).

Selecting a Region

When selecting a region, consider the following factors:

- Location

It is recommended that you select the closest region for lower network latency and quick access.

- If your target users are in Asia Pacific (excluding the Chinese mainland), select the **CN-Hong Kong**, **AP-Bangkok**, or **AP-Singapore** region.
- If your target users are in Africa, select the **AF-Johannesburg** region.
- If your target users are in Latin America, select the **LA-Santiago** region.

 **NOTE**

The **LA-Santiago** region is located in Chile.

- Resource price

Resource prices may vary in different regions. For details, see [Product Pricing Details](#).

Selecting an AZ

When deploying resources, consider your applications' requirements on disaster recovery (DR) and network latency.

- For high DR capability, deploy resources in different AZs within the same region.
- For lower network latency, deploy resources in the same AZ.

Regions and Endpoints

Before you use an API to call resources, specify its region and endpoint. For more details, see [Regions and Endpoints](#).

7 Working with Other Services

Figure 7-1 illustrates how an enterprise switch works with other Huawei Cloud services.

Figure 7-1 Interactions between an enterprise switch and other cloud services




Table 7-1 Interactions between an enterprise switch and other cloud services

Service	Interaction
Virtual Private Cloud (VPC)	VPCs can use enterprise switches to communicate with on-premises data centers at Layer 2.
Direct Connect	Direct Connect or VPN allows the communication between on-premises data centers and VPCs at Layer 3. Based on the Layer 3 network, enterprise switches establish a Layer 2 network between on-premises data centers and VPCs.
Virtual Private Network (VPN)	
Identity and Access Management (IAM)	On IAM, you can assign different permissions to different users to control their access to enterprise switch resources.

8 Billing

Billing Items

You are billed based on the enterprise switch type. Currently, basic, standard, and enhanced enterprise switches are supported.

For pricing details, see [Enterprise Switch Price Calculator](#).

Billing Mode

Enterprise switches can be billed on a yearly/monthly or pay-per-use basis. For billing details, see [Table 8-1](#).

Table 8-1 Enterprise switch billing details

Billing Mode	Description	Impact of Operations on Billing
Yearly/ Monthly	<p>If you buy a yearly/monthly enterprise switch, you need to make a one-off payment for the enterprise switch in your selected period. The price varies by enterprise switch specifications and is subject to the one on the payment page. Enterprise switches of the following specifications can be billed on a yearly/monthly basis:</p> <ul style="list-style-type: none">• Basic• Standard• Enhanced <p>NOTE The yearly/monthly billing mode is available only in regions LA-Santiago and AF-Johannesburg.</p>	<p>You can unsubscribe from a yearly/monthly subscription. Your actual usage fee and some preferential fees will be deducted from the refund amount.</p>

Billing Mode	Description	Impact of Operations on Billing
Pay-per-use	Billing starts immediately after an enterprise switch is created. Your enterprise switch is billed by the second, and settled by the hour. If the usage is less than an hour, you are billed based on the actual duration consumed. Enterprise switches of the following specifications can be billed on a pay-per-use basis: <ul style="list-style-type: none">• Basic• Standard• Enhanced	-

 **NOTE**

Currently, the following enterprise switch specifications are supported:

- Basic
 - Maximum Bandwidth: 3 Gbit/s
 - Maximum PPS: 500,000
 - Connected Subnets: 1
- Standard
 - Maximum Bandwidth: 5 Gbit/s
 - Maximum PPS: 1,000,000
 - Connected Subnets: 3
- Enhanced
 - Maximum Bandwidth: 10 Gbit/s
 - Maximum PPS: 2,000,000
 - Connected Subnets: 6

How Do I Renew an Enterprise Switch? What Will Happen If My Account Is in Arrears?

If your account is in arrears, you will be impacted for using your resources.

- If your yearly/monthly resource has expired and is not renewed, the resource enters the grace period. If you do not renew the monthly/yearly resource within the grace period, the resource enters a retention period after the grace period has expired.

You cannot perform any operations on yearly/monthly resources that are in the grace or retention period. To ensure that your services are not affected, **renew the resources** before they expire.

- If your pay-per-use resource is in arrears, the resource enters the grace period. If you do not pay off the arrears of the pay-per-use resource within the grace period, the resource enters a retention period after the grace period has expired.

You can still perform operations on pay-per-use resources in the grace period. However, you cannot perform any operations on them if they enter the retention period.