Cloud Load Balancer

Purchase Guide

Product Documentation

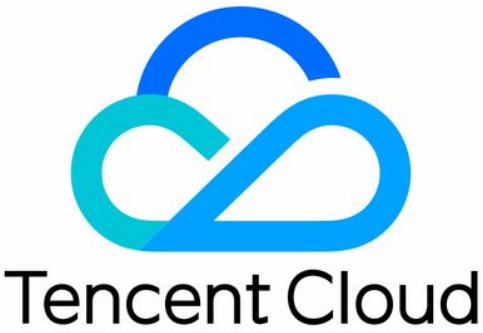

Contents

Purchase Guide
Billing Overview

Last updated：2024-01-04 09:44:16

This document describes CLB billing modes, billable items, billing cycles, and prices.

Note:

There are two types of Tencent Cloud accounts: bill-by-IP account and bill-by-CVM account. All Tencent Cloud accounts registered after 00:00:00 on June 17, 2020 are bill-by-IP accounts. If you have registered an account before that time, you can check your account type in the console as instructed in Checking Account Type. For the differences between these two account types, see Account Types.

Physical dedicated instances are deployed on dedicated clusters. To add or deactivate dedicated clusters, please contact your sales rep. The minimum service period of dedicated clusters is three months.

Billable items

The CLB service fees involve four parts, namely the instance fee, network fee, LCU fee and cross-region binding fee.

Bill-by-IP account

Bill-by-CVM account

Instance type	Network type	Instance fee	Public network fee	LCU fee	Cross-region binding fee
LCU-supported	Public network	Yes	Yes	Yes	Optional
	Private network		No		
Shared	Public network	Yes	Yes	No	Optional
	Private network		No		

Instance type	Network type	Instance fee	Public network fee	LCU Fee	Cross-region binding fee

LCU-supported	Public network	Yes	N/A (billed together with the CVM)	Yes	Optional
	Private network				
Shared	Public network	Yes	N/A (billed together with the CVM)	No	Optional
	Private network				

Note：

For bill-by-CVM accounts, when the client accesses IPv4 and Ipv6 NAT 64 CLB instances, the public network fee is billed on the CVM but not the CLB. For more details, see CVM Public Network Fee. If you have not purchased public network for you CVM, the client cannot access the CVM via a public network load balancer.

For bill-by-CVM accounts, when a client accesses an IPv6 CLB instance, the IPv6 public network bandwidth fee is counted into the bandwidth package (IPv4 and IPv6 NAT64 CLBs are not affected). When you create an IPv6 CLB instance, you need to specify the bandwidth cap and set the billing mode to Bandwidth Package. For more pricing details, see Bandwidth Package - Regular BGP Bandwidth Package.

Billing modes

Public network CLBs and private network CLBs support pay-as-you-go.

Bill-by-IP account

Bill-by-CVM account

Instance billing mode	Network billing mode	Billing overview
Pay-as-you-go	Traffic-based billing	Bill by the outgoing data volume
	Bandwidth packages	Bill by the higher one between inbound bandwidth peak and outbound bandwidth peak

Instance Billing Method	Network Billing Mode
Pay-as-you-go	N/A (Purchased together with the CVM)

Pricing

For the billable items of CLB, please refer to the documents below.

Instance fee

Network fee

LCU fee

Billing
Instance Pricing

Last updated：2024-01-04 09:44:16

An CLB instance incurs the instance fee even when it is idle (i.e., there are no access requests and no real server is bound).

Note:

Staring from 00:00:00, November 2, 2021 (UTC +8), the guaranteed performance of each shared CLB instance has been increased to 50,000 concurrent connections, 5,000 new connections per second, and 5,000 QPS. After the upgrade, the CLB instance fee is adjusted to 0.686 USD/day for most regions (1.029 USD/day for some regions). For more details, see CLB Price Changes.

Pay-as-You-Go

Billing cycle

The instance fee is calculated and billed on a daily basis.

Billing starts when the CLB instance is created successfully and ends when the termination task is initiated.

The minimum billable period is one hour.

Note:

When you create a pay-as-you-go CLB instance, an hour of instance fee (1/24 of the daily unit price) is deducted from your account balance. Make sure that your account balance is sufficient.

Pricing

Region	Instance fee (USD/hour)
Seoul, Mumbai, Virginia	0.0214
Tokyo	0.0231
Singapore, Jakarta, Silicon Valley	0.0239
Frankfurt, Hong Kong (China)	0.0257
Guangzhou, Shanghai, Nanjing, Beijing, Chengdu, Chongqing	0.0286
São Paulo	0.0323
Bangkok, Shanghai ADC	0.0429

See also

Billing Overview

Network Pricing

Last updated：2024-01-04 09:44:16

Only public network CLBs incur network fees. Private network CLB instances do not incur network fees.

Pay-as-You-Go

For pay-as-you-go CLB instances, the network billing modes include traffic-based billing and bandwidth package.

Pricing

Traffic-based billing

The volume of data (in GB) transferred over the public network is calculated and billed on an hourly basis. This mode is applicable to scenarios where the business traffic fluctuates greatly. It's recommended to use this mode if your bandwidth utilization is below 10%.

Note:

The billable traffic volume is calculated based on the traffic going from the CLB instance to the public network.

To prevent high bills due to traffic surges, you can set a bandwidth cap. When the bandwidth cap is reached, the exceeding packets are dropped and no fees are incurred.

The traffic units are 1024-based, which means 1 TB = 1,024 GB, and 1 GB = 1,024 MB.

Region	Public network fee (USD/GB)
Chinese mainland, Seoul, Hong Kong (China), Jakarta	0.12
Singapore	0.081
Frankfurt, Silicon Valley	0.077
Tokyo	0.13
Virginia	0.075
Bangkok, Mumbai	0.1
São Paulo	0.15

Bandwidth package

This billing mode bills multiple IPs in an aggregated manner and is applicable to large-scale businesses where public network instances have traffic peaks at different times. For more details about the pricing, see BGP Bandwidth Package.

See also

Billing Overview

LCU Fees

Last updated：2024-01-04 09:44:16

A Loadbalancer Capacity Unit (LCU) measures the performance metrics when CLB processes traffic. This document describes how the LCU charge is calculated.

Pay-as-You-Go

Note:

Only LCU-supported CLB instances will incur LCU usage fees.

Specification	Concurrent Connections per Minute	New Connections per Second (CPS)	Queries per Second (QPS)	Max Bandwidth
Standard	100,000	10,000	10,000	2Gbps
Advanced 1	200,000	20,000	20,000	4Gbps
Advanced 2	500,000	50,000	30,000	6Gbps
Super Large 1	1,000,000	100,000	50,000	10Gbps
Super Large 2	2,000,000	200,000	100,000	20Gbps
Super Large 3	5,000,000	500,000	200,000	40Gbps
Super Large 4	10,000,000	1,000,000	300,000	60Gbps

LCU metrics

In the pay-as-you-go mode, LCU usage is measured by using the following four metrics: CPS, concurrent connections per minute, processed traffic, and rule evaluations. Rule evaluations are supported only for HTTP and HTTPS protocols. The performance of a single LCU is related to the protocol type of the listener.

HTTP/HTTPS

Metric	Description	LCU Coefficient	Number of LCUs per Hour
CPS	Number of new connections per	25	Number of LCUs = Average

	second.		number of new connections per second within an hour ÷ LCU coefficient
Concurrent connections	Number of concurrent connections per minute.	3,000	Number of LCUs = Average number of concurrent connections per minute within an hour ÷ LCU coefficient
Processed traffic	The total amount of outbound and inbound traffic processed by CLB instances per hour, in GB.	1 GB	Number of LCUs = Total amount of traffic within an hour ÷ LCU coefficient
Rule evaluations (only for HTTP/HTTPS)	Product of QPS and the number of layer-7 rules processed by the CLB instances. The first 10 processed rules are free of charge: For over 10 processed rules: Number of rule evaluations = QPS × (Number of processed rules − 10). For 10 or fewer processed rules, Number of rule evaluations = QPS.	1,000	Number of LCUs = Average number of rule evaluations per second within an hour ÷ LCU coefficient

TCP/TCP SSL

Metric	Description	LCU Coefficient	Number of LCUs per Hour
CPS	Number of new connections per second.	800	Number of LCUs = Average number of new connections per second within an hour ÷ LCU coefficient
Concurrent connections per minute	Number of concurrent connections per minute.	100,000	Number of LCUs = Average number of concurrent connections per minute within an hour ÷ LCU coefficient
Processed traffic	The total amount of outbound and inbound traffic processed by CLB instances per hour, in GB.	1 GB	Number of LCUs = Total amount of traffic within an hour ÷ LCU coefficient

UDP/QUIC

Metric	Description	LCU	Number of LCUs per Hour

		Coefficient	
CPS	Number of new connections per second.	400	Number of LCUs = Average number of new connections per second within an hour ÷ LCU coefficient
Concurrent connections per minute	Number of concurrent connections per minute.	50,000	Number of LCUs = Average number of concurrent connections per minute within an hour ÷ LCU coefficient
Processed traffic	The total amount of outbound and inbound traffic processed by CLB instances per hour, in GB.	1 GB	Number of LCUs = Total amount of traffic within an hour ÷ LCU coefficient

In pay-as-you-go mode, the billing cycle of LCUs is one hour. The number of LCUs consumed in one hour is calculated based on the aforementioned metrics, and fees are charged based on the maximum number of LCUs consumed by the aforementioned metrics. The LCU fees for a single CLB instance is the sum of the LCU costs of the listeners for different protocols.

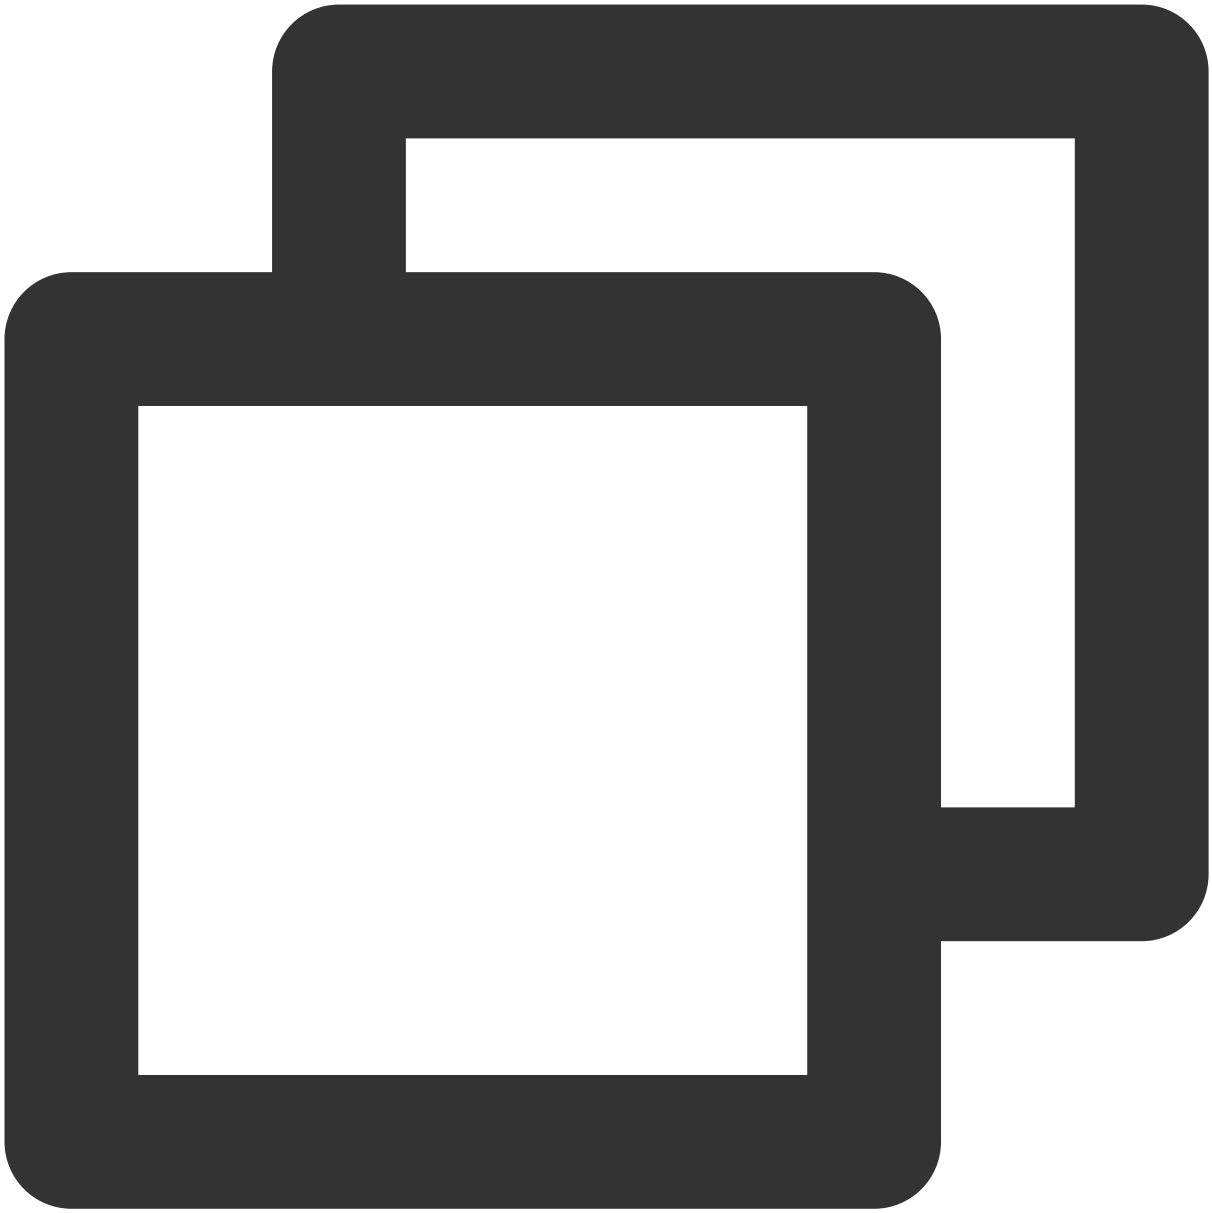

```
LCU fees per hour = MAX[Number of LCUs for concurrent connections, Number of LCUs f
```

Pricing

Region	LCU Unit Price (USD/Hour)	LCU Unit Price for New Purchases (USD/Hour)
All regions	0.0072	0.0059

Note:

Starting from 00:00:00, June 1, 2023 (UTC+8):
For new purchases of CLB pay-as-you-go instances, the LCU unit price is adjusted to 0.0059 USD/hour.
For existing CLB instances, the LCU fees and other billable items are not affected.
This price adjustment is always valid unless otherwise declared.

Billing examples

Example 1: HTTP/HTTPS
Example 2: TCP/UDP
Assume that your LCU-supported CLB instance receives an average of 100 new connections per second, each lasting 3 minutes, that a client sends an average of 400 queries per second and CLB processes a total of 1,000 KB of query and response data per second, and that you have configured an HTTP listener on port 80 and an HTTPS listener on port 443, with 20 forwarding rules to route your client queries. You are charged on the following metrics:

Metric Name	Use Case	LCU Coefficient	LCUs Quantity Calculation
New connections (per second)	In this example, your CLB instance receives 100 new connections per second.	25	Number of LCUs = 100 ÷ LCU Coefficient = 4
Concurrent connections (per minute)	In this example, your CLB instance receives 100 new connections per second, which means 6,000 new connections per minute. As each connection lasts 3 minutes, the maximum concurrent connections per minute = 6,000 × 3 = 18,000.	3,000	Number of LCUs = 18,000 ÷ LCU Coefficient = 6
Processed traffic (per hour)	In this example, your CLB instance processes 1,000 KB of data per second. This translates to 3.6 GB per hour.	1	Number of LCUs = 3.6 ÷ LCU Coefficient = 3.6
Rule evaluations (per second)	In this example, your CLB instance receives 400 queries per second on average. Therefore, 20 forwarding rules result in 4,000 rule evaluations per second for each query: (20 forwarding rules − 10 free rules) × 400.	1,000	Number of LCUs = 4,000 ÷ LCU Coefficient = 4

In this example, the maximum number of LCUs is consumed by the concurrent connections metric, which consumes 6 LCUs.
Therefore, hourly fees of LCUs = 0.0072 USD/LCU × 6 LCUs = 0.0432 USD, and monthly fees = 0.0432 USD/h × 24 hours × 30 days = 31.104 USD.

Assume that your CLB instance receives 100 new TCP connections (each transferring 1,000 processed bytes and each lasting 3 minutes) and 100 UDP flows (each transferring 1,000 processed bytes and each lasting 2 minutes). Your LCU fees will be calculated as follows:

TCP traffic

Metric Name	Use Cases	LCU Coefficient	LCUs Quantity Calculation
New connections (per second)	In this example, your CLB instance uses 100 new connections per second.	800	Number of LCUs = 100 ÷ LCU Coefficient = 0.125
Concurrent connections (per minute)	In this example, your CLB instance receives 100 new TCP connections per second, each lasting 3 minutes. Therefore, the maximum number of concurrent connections per minute = 6,000 × 3 = 18,000.	100,000	Number of LCUs = 18,000 ÷ LCU Coefficient = 0.18
Processed traffic (per hour)	In this example, your CLB instance transfers 1,000 processed bytes for each TCP client connection. Therefore, the amount of traffic processed per hour = 100 connections/s × 60s × 60 min × 1,000 bytes = 0.36 GB.	1	Number of LCUs = 0.36 ÷ LCU Coefficient = 0.36

UDP traffic

Metric Name	Use Case	LCU Coefficient	LCUs Quantity Calculation
New connections (per second)	In this example, your CLB instance uses 100 new UDP connections per second.	400	Number of LCUs = 100 ÷ LCU Coefficient = 0.125
Concurrent connections (per minute)	In this example, your CLB instance receives 100 new TCP connections per second, each lasting 2 minutes. Therefore, the maximum number of concurrent connections per minute = 100 × 60 × 2 = 12,000.	50,000	Number of LCUs = 12,000 ÷ LCU Coefficient = 0.24

Processed traffic (per hour)	In this example, your CLB instance transfers 1,000 processed bytes for each UDP client connection. Therefore, the amount of traffic processed per hour = 100 connections/s × 60s × 60 min × 1,000 bytes = 0.36 GB.	1	Number of LCUs = 0.36 ÷ LCU Coefficient = 0.36

The hourly fee for each protocol is calculated based on the maximum number of LCUs used by these metrics.
In the example of TCP traffic, the maximum number of LCUs is consumed by the processed traffic metric, which consumes 0.36 LCUs. Therefore, the hourly LCU fee for TCP = 0.0072 USD/LCU × 0.36 LCU = 0.002592 USD, and the monthly fee = 0.002592 USD/h × 24 hours × 30 days = 1.86624 USD.
In the example of UDP traffic, the maximum number of LCUs is consumed by the processed traffic metric, which consumes 0.36 LCU. Therefore, the hourly LCU fee for UDP = 0.0072 USD/LCU × 0.36 LCUs = 0.002592 USD, and the monthly fee = 0.002592 USD/h × 24 hours × 30 days = 1.86624 USD.
Then, the total hourly LCU fee of your LCU-supported CLB instance is 0.005184 USD/hour (with 0.002592 USD for TCP traffic and 0.002592 USD for UDP traffic).
The estimated total monthly LCU fee on average is 3.73248 USD/month (with 1.86624 USD for TCP traffic and 1.86624 USD for UDP traffic)

References

Billing Overview
Purchase Methods
Upgrading CLB Instance

Cross-region Binding Fee

Last updated：2024-01-04 10:35:56

Public Network Cloud Load Balancer, once the cross-region binding 1.0 is enabled, will incur cross-region fees. No fees will be charged when it is disabled. Private Network Cloud Load Balancer does not support this function. If you utilize cross-region binding 2.0 function, the cross-region fee is computed on Cloud Connect Network, but not on CLB. If this function is not configured, there is no need for concern. This document describes the billing details of cross-region binding 1.0.

Note:

If you require the use of the cross-region binding feature, please submit a ticket to consult.

Billing Mode

Binding CLB to CVM in the same region: Free.

Binding CLB to CVM in different regions 1.0: The billing is based on peak value of the daily inbound and outbound bandwidth after 95 peak shaving.

Binding CLB to CVM in different regions 2.0: Billing through Cloud Connect Network.

Cross-Region Binding 1.0 calculation formula

Daily expense = Peak of the daily inbound and outbound bandwidth after 95% of peak shaving × Unit price at the bandwidth tier

Peak value of the daily outbound and inbound bandwidth after 95 peak shaving : A collection is performed every 5 minutes. Points with a bandwidth greater than 0 are valid. All valid points are sorted from high to low. The top 5% points are removed, and the remaining highest point is the peak value after 95 peak shaving (denoted as Max95), which is the billing bandwidth.

Tiered unit price: The unit price of the tier in which Max95 is located.

Cross-Region Binding 1.0 Billing Price

Please refer to the table below for the unit price of the tier of bandwidth:

Feature	Bandwidth Range (Mbps)	Price (Unit: USD/Mbps/Day)	
		Intercommunication in Mainland China (Excluding Hong Kong	Other Regions

		(China), Macao (China), and Taiwan (China))	
Cross-Region binding	(0 , 20)	3.19	For price inquiries, please consult with business affairs.
	[20 ,100)	1.98	
	[100 , 500)	For price inquiries, please consult with business affairs.	
	[500 , 2000)		
	≥ 2000		
Binding in the same region	Free		

Note:

Inbound and outbound bandwidth ≤ 1Kbps is not charged.

Cross-Region Binding 1.0 Billing Example

If Cloud Load Balancer is located in Shanghai and the backend host is in Guangzhou, the peak value of daily outbound bandwidth after peak shaving is 20 Mbps and the peak value of daily inbound bandwidth after peak shaving is 30 Mbps, then:

Peak value of the daily outbound and inbound bandwidth after 95 peak shaving: 30 Mbps.

Tiered unit price for 30Mbps: 1.98 USD/Mbps/day.

Daily cost = Peak value of the daily outbound and inbound bandwidth after 95 peak shaving × tiered unit price = 30 × 1.98 = 59.4 USD, which is charged by the Cloud Load Balancer.

Purchase Methods

Last updated：2024-01-04 09:44:16

Tencent Cloud allows you to purchase a CLB instance on the official purchase page or via API. The two methods are detailed below:

Purchasing a CLB Instance on the Official Purchase Page

You can purchase a CLB instance on Tencent Cloud official website. There are two types of Tencent Cloud accounts, namely the bill-by-IP account and bill-by-CVM account. The accounts created after 00:00:00 on June 17, 2020 are of the bill-by-IP type. If you have created your account before that time, you can check your account type in the console as instructed in Checking Account Type.

1. Log in to the Tencent Cloud console and go to the CLB purchase page.
2. Select the following CLB configuration items as needed:

Bill-by-IP account

Parameter	Description
Billing mode	Supports pay-as-you-go billing.
Region	Select a region. For more information on the regions supported by CLB, see Region List.
Instance type	Supports the CLB instance type only. Starting from October 20, 2021, classic CLB instances can no longer be purchased. For more information, see Classic CLB End-of-Sale Notice.
Network type	Supports two network types: public network and private network. For more information, see Network Types. Public network: CLB is used to distribute requests from a public network. Private network: CLB is used to distribute requests from the Tencent Cloud private network. A private network instance does not support the following configuration items, and therefore they are not displayed by default: EIP, IP version, ISP, instance specification, network billing mode, and bandwidth cap. The supported network types vary by billing mode: In pay-as-you-go billing mode, both the public and private network types are supported.
EIP	If EIP is not selected, Tencent Cloud will assign you a public network CLB instance whose public IP address cannot be changed. If EIP is selected, Tencent Cloud will assign you an EIP and a private network CLB instance, which has the similar features of public network CLB. (Only pay-as-you-go public network CLB instances allow you to select an EIP.)

	This feature is currently in beta. To try it out, submit a ticket. For the use limits, see Use Limits.
IP version	Supports the following CLB IP versions: IPv4, IPv6, and IPv6 NAT64. Only pay-as-you-go instances support the IPv6 version. For more information about other restrictions, see IP Versions. IPv6 CLB is currently in beta. To try it out, submit a ticket.
Network	CLB supports classic network and VPC. The classic network is a public network resource pool for all Tencent Cloud users. The private IPs of all CVMs are assigned by Tencent Cloud. You cannot customize IP ranges or IP addresses. A VPC is a logically isolated network space in Tencent Cloud. In a VPC, you can customize IP ranges, IP addresses, and routing policies. A VPC is more suitable for use cases requiring custom configurations. Besides, the overall classic network products were officially discontinued on December 31, 2022. For details, see Ending Support for Classic Network. We recommend you choose a VPC.
ISP	Supports the following ISP types: BGP (multi-line), China Mobile, China Telecom, and China Unicom. In pay-as-you-go billing mode, all of the above four options are supported. Currently, the static single-line IP is supported only in Guangzhou, Shanghai, Nanjing, Jinan, Hangzhou, Fuzhou, Beijing, Shijiazhuang, Wuhan, Changsha, Chengdu, and Chongqing. For the support information in other regions, see the console. If you want to try it out, contact the sales rep for application. Once your application is approved, you can select an ISP (China Mobile, China Unicom, or China Telecom) on the purchase page.
Primary/Secondary availability zone	The primary availability zone (AZ) is an AZ that currently sustains the traffic. The secondary AZ does not sustain traffic by default and will be used only when the primary AZ is unavailable. Currently, only IPv4 CLB instances in the Guangzhou, Shanghai, Nanjing, Beijing, Chengdu, Shenzhen Finance, Hong Kong (China), Seoul, Frankfurt, Singapore regions support primary/secondary AZs.
Instance specification	Supports shared and LCU-supported instances. Shared instances guarantee performance according to their specifications. A single instance can sustain up to 50,000 concurrent connections, 5,000 new connections per second, and 5,000 queries per second. An LCU-supported instances provide guaranteed performance according to specifications. A single instance can sustain up to Ten million concurrent connections, 1 million new connections per second, and 300,000 queries per second.
Network billing mode	Supports the following network billing modes: bill-by-bandwidth (monthly subscription and hourly bandwidth), bill-by-traffic, and bandwidth package. A pay-as-you-go instance supports three network billing modes: bill-by-bandwidth (hourly bandwidth), bill-by-traffic, and bandwidth package.
Bandwidth cap	The public network bandwidth cap of a shared CLB is 2 Gbps. For the private network bandwidth cap, it's suggested to set a value within 5 Gbps.

	The bandwidth cap of an LCU-supported CLB instance depends on the selected specification. For details, see Instance Specifications Comparison.
Project	Select a project.
Tag	Select a tag key and value. You can also create a tag as instructed in Creating Tags and Binding Resources.
Instance name	The name can contain up to 60 characters, including letters, numbers, hyphens, underscores, and dots. If it is not specified, a name will be automatically generated by default.

Bill-by-CVM account

Parameter	Description
Billing mode	Supports pay-as-you-go billing only.
Region	Select a region. For more information on the regions supported by CLB, see Region List.
Instance type	Supports the CLB instance type only. Starting from October 20, 2021, classic CLB instances can no longer be purchased. For more information, see Classic CLB End-of-Sale Notice.
Network type	Supports two network types: public network and private network. For more information, see Network Types. Public network: CLB is used to distribute requests from a public network. Private network: CLB is used to distribute requests from the Tencent Cloud private network. A private network instance does not support the following configuration items, and therefore they are not displayed by default: IP version, ISP, and instance specification.
IP version	Supports the following CLB IP versions: IPv4, IPv6, and IPv6 NAT64. For more information about use limits, see IP Versions. IPv6 CLB is currently in beta. To try it out, submit a ticket.
Network	CLB supports classic network and VPC. The classic network is a public network resource pool for all Tencent Cloud users. The private IPs of all CVMs are assigned by Tencent Cloud. You cannot customize IP ranges or IP addresses. A VPC is a logically isolated network space in Tencent Cloud. In a VPC, you can customize IP ranges, IP addresses, and routing policies. A VPC is more suitable for use cases requiring custom configurations. Besides, the overall classic network products were officially discontinued on December 31, 2022. For details, see Ending Support for Classic Network. We recommend you choose a VPC.
ISP	Supports the following ISP types: BGP (multi-line), China Mobile, China Telecom, and China Unicom.

	Currently, the static single-line IP is supported only in Guangzhou, Shanghai, Nanjing, Jinan, Hangzhou, Fuzhou, Beijing, Shijiazhuang, Wuhan, Changsha, Chengdu, and Chongqing. This feature is in beta. To try it out, submit a ticket. For the support information in other regions, see the console. If you want to try this feature, contact the sales rep for application. Once your application is approved, you can select an ISP (China Mobile, China Unicom, or China Telecom) on the purchase page.
Instance specification	Supports shared and LCU-supported instances. Shared instances guarantee performance according to their specifications. A single instance can sustain up to 50,000 concurrent connections, 5,000 new connections per second, and 5,000 queries per second. An LCU-supported instances provide guaranteed performance according to specifications. A single instance can sustain up to Ten million concurrent connections, 1 million new connections per second, and 300,000 queries per second.
Project	Select a project.
Tag	Select a tag key and value. You can also create a tag as instructed in Creating Tags and Binding Resources.
Instance name	The name can contain up to 60 characters, including letters, numbers, hyphens, underscores, and dots. If it is not specified, a name will be automatically generated by default.

3. After completing the above configuration, confirm the quantity and fees and click **Buy now**.

Pay-as-you-go billing mode: In the **Confirm** pop-up window, click **OK**.

4. After successful purchase, CLB will be activated and you can configure and use the CLB instance.

Purchasing a shared instance

1. Log in to the Tencent Cloud console and go to the CLB purchase page.

2. Set the shared instance configuration items by referring to the steps in Purchasing a CLB Instance on the Official Purchase Page and select **Shared** for **Instance specification**.

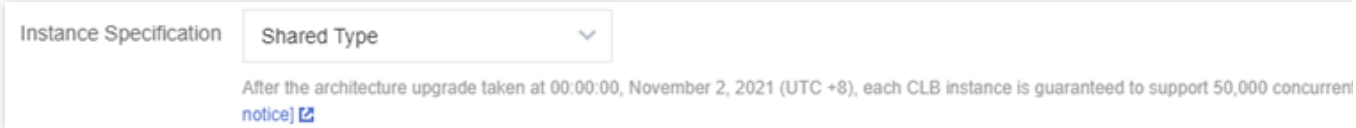

3. Complete the subsequent operations by referring to the steps in Purchasing a CLB Instance on the Official Purchase Page.

Purchasing an LCU-supported instance

1. Log in to the Tencent Cloud console and go to the CLB purchase page.

2. Set the LCU-supported instance configuration items by referring to the steps in Purchasing a CLB Instance on the Official Purchase Page and select **LCU-supported** for **Instance specification**.

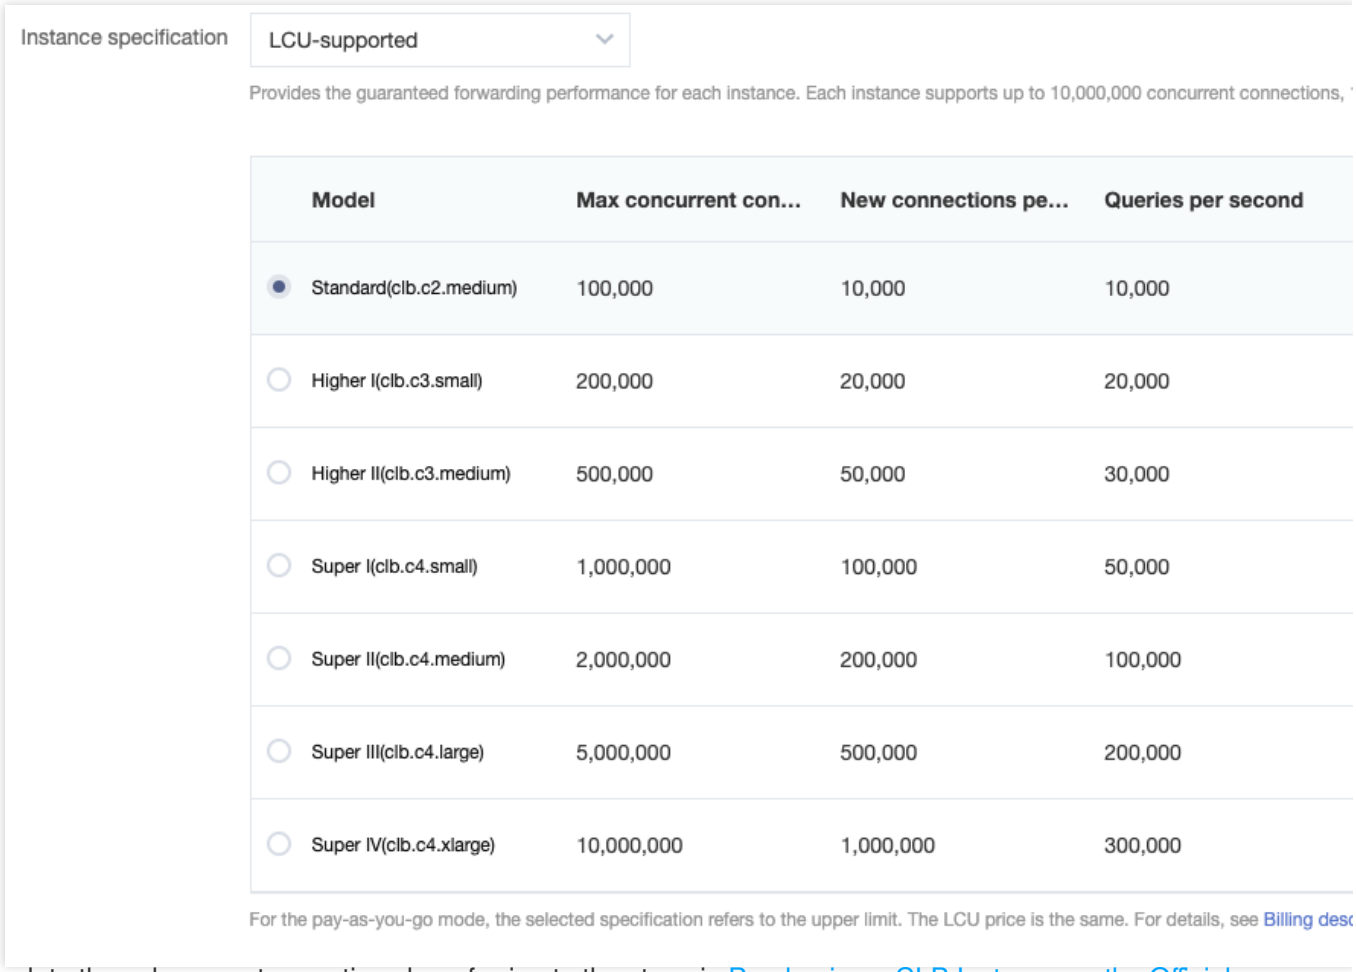

3. Complete the subsequent operations by referring to the steps in Purchasing a CLB Instance on the Official Purchase Page.

Purchasing a CLB Instance via an API

To purchase a CLB instance via an API, see CreateLoadBalancer.

Related Operations

You can create a listener for a CLB instance as instructed in CLB Listener Overview.
You can bind a CLB listener to a real server as instructed in Real Server Overview.

References

Product Attribution Selection

Payment Overdue

Last updated：2024-01-04 09:44:16

Note:

If you are a customer of a Tencent Cloud partner, the rules regarding resources when there are overdue payments are subject to the agreement between you and the partner.

Pay-as-You-Go

Note:

Delete unnecessary pay-as-you-go resources to stop billing.

Since your actual resource consumption is constantly changing, some slight discrepancies may exist for your stated balance.

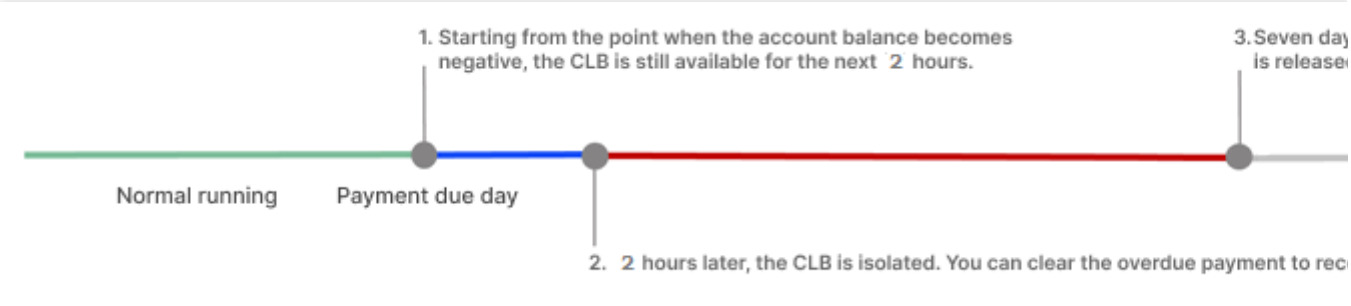

Balance Alert

The system estimates the time that your account balance is about to run out based on the current balance and your usage in the past 24 hours. If the estimated time is within 5 days, an alert is sent to the Tencent Cloud account owner and collaborators subscribed to the alert in the specified method (email, SMS, and the Message Center).

Overdue Payment Alert

Pay-as-you-go resources are billed on an hourly basis. When the account balance becomes negative (Point 1 in the figure above), an alert is sent to the Tencent Cloud account owner and collaborators subscribed to the alert in the specified method (email, SMS, and the Message Center).

Overdue Policies

Starting from the point when your account balance becomes negative, the pay-as-you-go load balancers are still available for the next 2 hours, and the billing continues. 2 hours later (Point 2 in the figure above), the load balancers

will be isolated and suspended. The billing stops.

Isolated and suspended CLBs will be released seven days later, and cannot be recovered.

Note:

A CLB instance will NOT be unbound from the CVM instance automatically. When a CVM instance is isolated, it **will NOT be unbound** from the CLB instance either.

Product Attribute Selection

Last updated：2024-01-04 17:46:37

On CLB purchase page, you can select CLB instances of different attributes. This document describes how to purchase instances based on business scenarios.

Regions

We recommend selecting the region closest to your client to reduce delay and accelerate the download speed.
CLB can only forward traffic to CVM instances in the same region. Please select the region where your CVM instance is located before creating a CLB instance.
CLB can forward traffic to CVM instances in multiple availability zones in the same region. For example, a CLB instance in Beijing can forward traffic to CVM instances in Beijing Zone 1, Beijing Zone 2, and Beijing Zone 3.
A public network CLB (formerly "application CLB") instance supports cross-region binding of CVM. You can select the region of real servers and bind them across VPCs or regions. For example, you can bind a CLB instance in Beijing to a CVM instance in Shanghai to forward the client traffic to the latter. For more information, please see CLB Instance Cross-Region Binding.

Instance Types

CLB instances can be classified into two types: CLB (formerly "application CLB") and classic CLB. CLB covers all features of classic CLB. Based on product features and performance, we recommend CLB. For a detailed comparison, please see Instance Types.

Network Types

Public network CLB

If you need to use CLB to distribute requests from the public network, please select "Public Network".
Public network CLB instances obtain requests from the client via the internet and distribute them to the bound real servers. After you create a public network CLB instance, Tencent Cloud will assign it a VIP address to be resolved by DNS server. Public network CLB allows users to add CNAME and A records, and map them to custom domain names. VIPs of public network CLB instances are fixed public IPs. They can receive HTTP, HTTPS, TCP, UDP, and other requests from the client.
Use Cases

A server cluster is used to provide services to public networks. A unified entry is required, while requests from users of public networks should be assigned to server clusters.
Users of different ISPs are connected to networks closest to them to accelerate network access.

Public VIP types

Regular IP: Regular BGP IP, which is suitable for users want to balance the network quality and cost.
Single-line IP: Accessing the public network via a single ISP, which is low-cost and convenient for self-scheduling.

Billing

For a bill-by-CVM account, only the instance fee is charged for purchasing a public network CLB instance. The service fees incurred by public network bandwidth and traffic are charged on its real server. For more information, please see Bill-by-CVM Account Billing Description.
For a bill-by-IP account, the instance fee and public network fee are charged for purchasing a public network CLB instance. For more information, please see Bill-by-EIP/CLB Account Billing Description.

Private network CLB

If you need to use CLB to distribute requests from private networks, please select "Private Network".
Both the client and server of private network CLB are in Tencent Cloud and can only be accessed within Tencent Cloud. Access via the internet is not allowed (no public network IP address). Tencent Cloud will assign a private IP to CLB from the network you have selected (such as a VPC). The request traffic will be routed to the backend CVM instance via this IP address.
If your application has multiple layers, such as web servers that can communicate over the internet or database servers that can only communicate over the private network, you can create a framework with both public and private network CLB instances. You can then connect all web servers to the public network CLB instance and database servers to the private network CLB instance. The public network CLB instance receives requests from the internet and sends them to the backend web servers. After being processed, requests to database will be sent to the private network CLB instance, which will then route the requests to database servers.

Use Cases

Tencent Cloud has more than one internal servers. Requests from the client need to be allocated to servers properly.
Fault tolerance and restoration are needed for the internal server cluster.
The service provider wants to block their own physical IP addresses and provide imperceptible services to the client.

Billing

The usage of private network CLB instances incur charges. But you don't need to pay the related public network fee.

IP Versions

The IP version of a CLB instance. Available versions: IPv4, IPv6, and IPv6 NAT64.

Enabling Anycast

If you enable Anycast during the CLB instance creation, an Anycast CLB instance will be created. Anycast CLB is a load balancing service that supports dynamic acceleration across regions. CLB VIP is published in multiple regions. The client connects to the nearest POP, and the traffic will be forwarded to a CVM instance through the high-speed internet of Tencent Cloud IDC. For more information, please see Creating an Anycast Instance.

Instance Name

CLB instance name, which can contain 1 to 60 characters.

Purchase Quantity

Number of purchased CLB instances.