

# Alibaba Cloud

## Server Load Balancer Product Introduction









Document Version: 20210202

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

# Table of Contents

1.What is CLB? .....	05
2.High availability .....	07
3.Architecture .....	10
4.Features .....	14
5.Scenarios .....	17
6.Terms .....	20

# 1. What is CLB?

This topic provides an overview of Classic Load Balancer (CLB). CLB distributes inbound network traffic across multiple Elastic Compute Service (ECS) instances that act as backend servers based on forwarding rules. You can use CLB to improve the responsiveness and availability of your applications.

## Overview

After you add ECS instances that are deployed in the same region to a CLB instance, CLB uses virtual IP addresses (VIPs) to virtualize these ECS instances into backend servers in a high-performance server pool that ensures high availability. Client requests are distributed to the ECS instances based on forwarding rules.

CLB checks the health status of the ECS instances and automatically removes unhealthy ones from the server pool to eliminate single points of failure (SPOFs). This enhances the resilience of your applications. You can also use CLB to defend your applications against distributed denial of service (DDoS) attacks.

## Components

CLB consists of three components:

- CLB instances


A CLB instance is a running CLB service entity that receives traffic and distributes traffic to backend servers. To get started with CLB, you must create a CLB instance and add at least one listener and two ECS instances to the CLB instance.

- Listeners

A listener checks client requests and forwards them to backend servers. It also performs health checks on backend servers.

- Backend servers

ECS instances are used as backend servers to receive distributed requests. You can separately add ECS instances to the server pool, or use vServer groups or primary/secondary server groups to add and manage ECS instances in batches.



## Benefits

- High availability

CLB is designed with full redundancy that avoids SPOFs and supports zone-disaster recovery.

CLB can be scaled based on application loads and can provide continuous service during traffic fluctuations.

- High scalability

You can increase or decrease the number of backend servers to adjust the load balancing capability of your applications.

- Cost-effectiveness

CLB can save 60% of load balancing costs compared with using traditional hardware solutions.

- Security

You can use CLB with Apsara Stack Security to defend your applications against up to 5 Gbit/s DDoS attacks.

- High concurrency

A CLB cluster supports hundreds of millions of concurrent connections and a single CLB instance supports tens of millions of concurrent connections.

## 2.High availability

This topic describes the high-availability architecture of CLB. You can use CLB in concert with DNS to implement geo-disaster recovery. CLB is designed to offer a multi-zone service availability of 99.99% and a single-zone service availability of 99.90%.

### High availability of the CLB architecture

CLB instances are deployed in clusters to synchronize sessions and protect backend servers from SPOFs, improving redundancy and ensuring service stability. Layer-4 CLB uses the open-source Linux Virtual Server (LVS) and Keepalived software to balance loads, whereas Layer-7 CLB uses Tengine. Tengine, a web server project launched by Taobao, is based on NGINX and adds advanced features dedicated for high-traffic websites.


Requests from the Internet reach an LVS cluster along Equal-Cost Multi Path (ECMP) routes. In the LVS cluster, each machine uses multicast packets to synchronize sessions with the other machines. At the same time, the LVS cluster performs health checks on the Tengine cluster and removes unhealthy machines from the Tengine cluster to ensure the availability of Layer-7 CLB.

Best practice:

You can use session synchronization to prevent persistent connections from being affected by server failures within a cluster. However, for short-lived connections or if the session synchronization rule is not triggered by the connection (the three-way handshake is not completed), server failures in the cluster may still affect user requests. To prevent session interruptions caused by server failures within the cluster, you can add a retry mechanism to the service logic to reduce the impact on user access.

### The high-availability solution with one CLB instance


To provide more stable and reliable load balancing services, you can deploy CLB instances across multiple zones in most regions to achieve cross-data-center disaster recovery. Specifically, you can deploy a CLB instance in two zones within the same region whereby one zone acts as the primary zone and the other acts as the secondary zone. If the primary zone suffers an outage, a failover is triggered to redirect requests to the servers in the secondary zone within approximately 30 seconds. After the primary zone is restored, traffic will be automatically switched back to the servers in the primary zone.

 **Note** Zone-disaster recovery is implemented between the primary and secondary zones. CLB implements failovers only when the whole CLB cluster within the primary zone is unavailable or fails, for example, due to power outage or optical cable failures. A failover will not be triggered when a single backend server fails.

Best practice:

1. We recommend that you create CLB instances in regions that support primary/secondary deployment for zone-disaster recovery.
2. You can choose the primary zone for your CLB instance based on the distribution of ECS instances. That is, select the zone where most of the ECS instances are located as the primary zone for minimized latency.

However, we recommend that you do not deploy all ECS instances in the primary zone. When you develop a failover solution, you must deploy several ECS instances in the secondary zone to ensure that requests can still be distributed to backend servers in the secondary zone for processing when the primary zone experiences a downtime.




## The high-availability solution with multiple CLB instances

In the context of one CLB instance, traffic distribution for your applications can still be compromised by network attacks or invalid CLB configurations, because the failover between the primary zone and the secondary zone is not triggered. As a result, the load-balancing performance is impacted. To avoid this situation, you can create multiple CLB instances to form a global load-balancing solution and achieve cross-region backup and disaster recovery. Also, you can use the instances with DNS to schedule requests so as to ensure service continuity.

Best practice:

You can deploy CLB instances and ECS instances in multiple zones within the same region or across different regions, and then use DNS to schedule requests.



## The high-availability solution with backend ECS instances

With health check enabled, CLB verifies the availability of backend ECS instances (or backend servers), and thus improves the availability of frontend services by minimizing downtime that is caused by health issues of ECS instances.



After you enable the health check feature, when an ECS instance is detected unhealthy, CLB distributes new requests to other healthy ECS instances. CLB will only send requests to this backend ECS instance when it is restored and considered healthy. For more information, see [Health check overview](#).

Best practice:

Make sure health check is enabled and properly configured. For more information, see [Configure health checks](#).

## 3. Architecture


This topic describes the CLB architecture. CLB instances are deployed in clusters to synchronize sessions and protect backend servers from SPOFs, improving redundancy and ensuring service stability. CLB supports Layer-4 load balancing of Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) traffic, and Layer-7 load balancing of HTTP and HTTPS traffic.

CLB forwards client requests to backend servers by using CLB clusters and receives responses from backend servers over internal networks.


### CLB design

CLB provides Layer-4 and Layer-7 load balancing.

- Layer-4 CLB combines the open-source Linux Virtual Server (LVS) with Keepalived to balance loads, and implements customized optimizations to meet cloud computing requirements.
- Layer-7 CLB uses Tengine to balance loads. Tengine is a web server project launched by Taobao. Based on NGINX, Tengine has a wide range of advanced features optimized for high-traffic websites.




Layer-4 CLB runs in a cluster of LVS machines, as shown in the following figure. This cluster deployment model strengthens the availability, stability, and scalability of the load balancing service in abnormal cases.



In an LVS cluster, each machine uses multicast packets to synchronize sessions with the other machines. As shown in the following figure, Session A established on LVS1 is synchronized to the other LVS machines after the client sends three data packets to the ECS instance. Solid lines indicate the current active connections, while dotted lines indicate that the session requests will be sent to other normally working machines if LVS1 fails or is being maintained. This way, you can perform hot upgrades, machine failure maintenance, and cluster maintenance without affecting the operation of business applications.

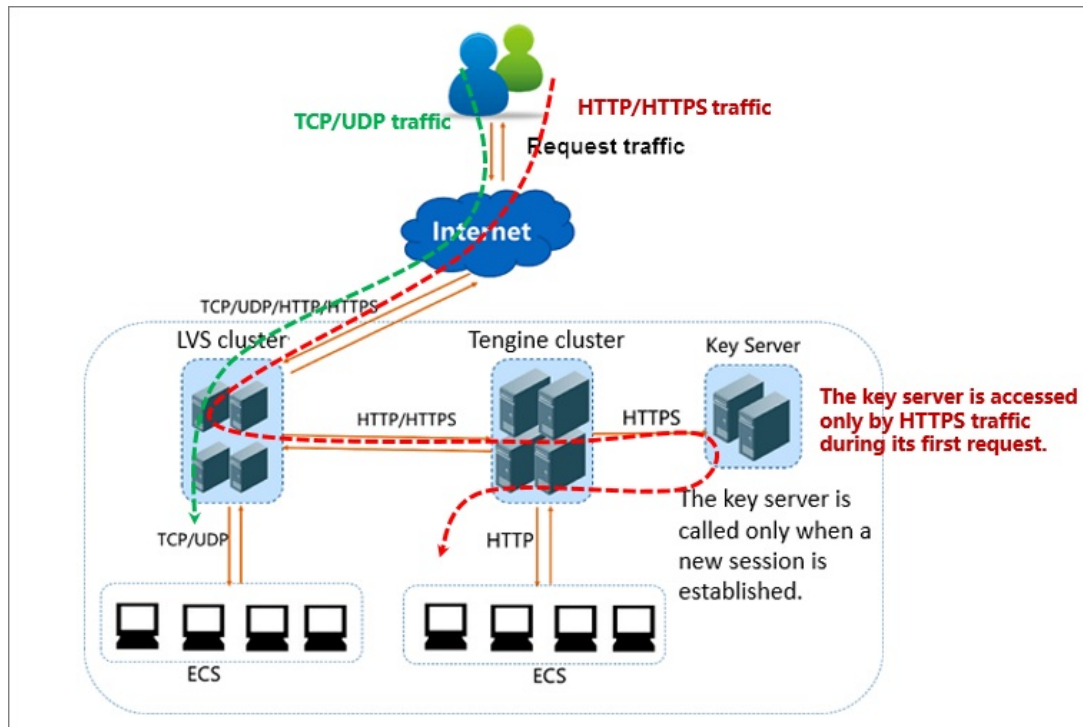
**Note** If a connection is not established (the three-way handshake is not completed), or if a connection has been established but session synchronization is not triggered during a hot upgrade, your service may be interrupted. In this case, you must re-initiate the connection from the client.



## Inbound network traffic flow

CLB distributes incoming traffic according to forwarding rules configured in the console or by using API operations. The following figure shows the inbound network traffic flow.

Inbound network traffic flow



1. For TCP, UDP, HTTP, and HTTPS protocols, the incoming traffic must be forwarded through the LVS cluster first.
2. The large amounts of inbound traffic is distributed evenly among all node servers in the LVS cluster, and the node servers synchronize sessions to ensure high availability.
  - o For Layer 4 listeners (either UDP or TCP), the node servers in the LVS cluster distribute requests directly to backend ECS instances according to the configured forwarding rules.
  - o For Layer 7 listeners that use the HTTP frontend protocol, the node servers in the LVS cluster first distribute requests to the Tengine cluster. Then, the node servers in the Tengine cluster distribute the requests to backend ECS instances according to the configured forwarding rules.
  - o For Layer 7 listeners that use the HTTPS frontend protocol, the request distribution is similar to the HTTP protocol except that the system calls the Key Server to validate certificates and decrypt data packets before requests are distributed to backend ECS instances.

## Outbound network traffic flow

CLB communicates with backend ECS instances by using the internal network.

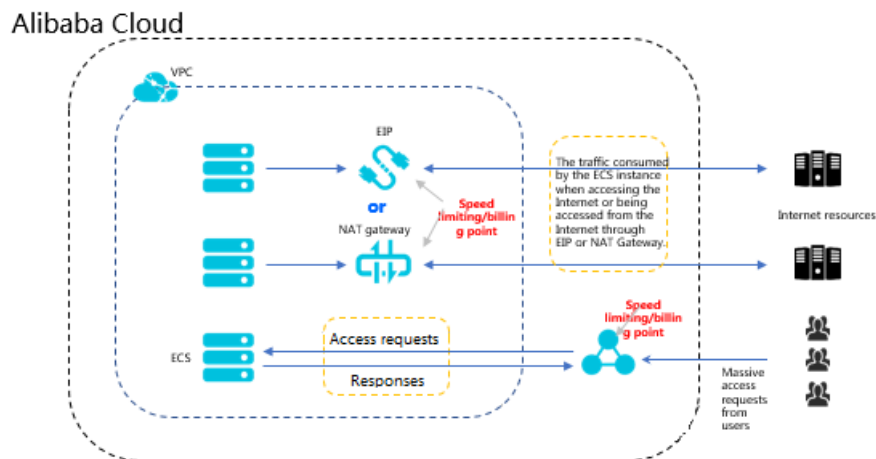
- If backend ECS instances handle the traffic distributed only from CLB, no public bandwidth (ECS instances, EIPs, NAT gateways, and public IP addresses) is required, and you do not need to purchase public bandwidth.

**Note** Previously created ECS instances are directly allocated with public IP addresses. You can view the public IP addresses by using the `ipconfig` command. If these ECS instances process requests only through CLB, no traffic fees are incurred for Internet traffic even if traffic statistics are read at the public network interface (NIC).

- If you want to provide external services through backend ECS instances, or backend ECS instances need to access the Internet, you must configure at least one of the following services: an ECS instance, a public IP address, an Elastic IP Address (EIP), or a NAT gateway.

The following figure shows the outbound network traffic flow.

Outbound network traffic flow




General principle: The traffic goes out from where it comes in.

- Outbound traffic from CLB instances is charged, and the speeds at which it is sent depend on the current network capacity. However, you are not charged for inbound traffic and internal communications between CLB instances and backend ECS instances.
- Traffic from an EIP or from a NAT gateway incurs charges. The speed at which EIP or NAT gateway traffic is sent depends on the current network capacity. If you configure public bandwidth when you purchase an ECS instance, traffic is sent at speeds dependent on the current network capacity, and is charged.
- CLB supports dynamic access to the Internet. Specifically, if a backend ECS instance needs to access the Internet, you must first configure public bandwidth for it, and purchase an EIP or a NAT gateway.
- The public bandwidth that you configured when you create ECS instances, EIPs, and NAT gateways allow ECS instances to access the Internet or be accessed from the Internet, but they cannot forward traffic or balance traffic loads.

## 4.Features

This topic describes the key features of Apsara Stack SLB and whether Layer-4 SLB and Layer-7 SLB support the features.

Feature	Layer-4 SLB	Layer-7 SLB
<p>Scheduling algorithms</p> <p>SLB supports the following scheduling algorithms: round-robin, weighted round-robin (WRR), weighted least connections (WLC), and consistent hash.</p>	✓	<p>✓</p> <div> <p> <b>Note</b></p> <p>Currently, Layer-7 SLB does not support the scheduling algorithm of consistent hash.</p> </div>
<p>Health check</p> <p>SLB checks the health status of backend servers. If a backend server is declared as unhealthy, the corresponding SLB instance stops distributing inbound traffic to it and distributes the traffic to other healthy backend servers.</p>	✓	✓
<p>Session persistence</p> <p>SLB provides session persistence. In a session, the SLB instance can distribute all of the requests from the same client to the same backend server.</p>	✓	✓
<p>Access control</p> <p>SLB uses whitelists and blacklists to control access to your applications.</p>	✓	✓
<p>High availability</p> <p>SLB can forward inbound traffic to backend servers in different zones. In addition, SLB is deployed in primary/secondary mode in most regions, meaning that an SLB instance is deployed across two zones within the same region. One zone is the primary zone and the other zone is the secondary zone. If the primary zone fails, the SLB instance will automatically fail over to the secondary zone, ensuring the high availability of the SLB service.</p>	✓	✓

Feature	Layer-4 SLB	Layer-7 SLB
<b>Security</b> SLB can be used with Apsara Stack Security to provide a DDoS protection capability of 5 Gbit/s.	✓	✓
<b>Public and internal network load balancing</b> SLB can load-balance both public network traffic and internal network traffic. You can create an internal SLB instance to balance traffic within a VPC, or create a public SLB instance to balance traffic coming from the public network.	✓	✓
<b>Monitoring</b> You can use CloudMonitor to view the monitoring data of SLB instances such as the number of connections and inbound and outbound traffic volumes.	✓	✓
<b>IPv6 support</b> SLB can forward requests from IPv6 clients.	✓	✓
<b>Health check logs</b> SLB stores health check logs of backend servers generated within three days by default. You can store all health check logs in OSS for troubleshooting.	✓	✓
<b>Domain name forwarding and URL forwarding</b> Layer-7 SLB supports configuration of domain name forwarding rules and URL forwarding rules to forward requests from different domain names or URLs to different backend servers.	-	✓
<b>Certificate management</b> SLB provides centralized certificate management for HTTPS listeners. You do not need to upload certificates to backend servers. Deciphering is performed on SLB instances to reduce the CPU utilization of backend servers.	-	✓
<b>SNI support</b> SLB allows you to associate multiple certificates to an HTTPS listener to distribute requests from different domain names to different backend servers.	-	✓
<b>Redirection</b> SLB supports redirecting HTTP requests to HTTPS.	-	✓

Feature	Layer-4 SLB	Layer-7 SLB
<p>WS/WSS support</p> <p>WebSocket is a new HTML5 protocol. It provides bidirectional communication channels between clients and servers, saving server resources and bandwidth, and achieving real-time communication.</p>	-	✓
<p>HTTP/2 support</p> <p>HTTP/2 is the second version of HTTP. It is backward compatible with HTTP1.X and significantly improves performance.</p>	—	✓



## 5.Scenarios

Classic Load Balancer (CLB) can be used to improve the availability and reliability of applications with high access traffic.

### Balance the loads of your applications

You can configure listening rules to distribute heavy traffic among ECS instances that are attached as backend servers to CLB instances. You can also use the session persistence feature to forward all of the requests from the same client to the same backend ECS instance to enhance access efficiency.

### Scale your applications

You can extend the service capability of your applications by adding or removing backend ECS instances to suit your business needs. CLB can be used for both web servers and application servers.

### Eliminate single points of failure (SPOFs)


You can attach multiple ECS instances to a CLB instance. When an ECS instance malfunctions, CLB automatically isolates this ECS instance and distributes inbound requests to other healthy ECS instances, ensuring that your applications continue to run properly.

### Implement zone-disaster recovery (multi-zone disaster recovery)

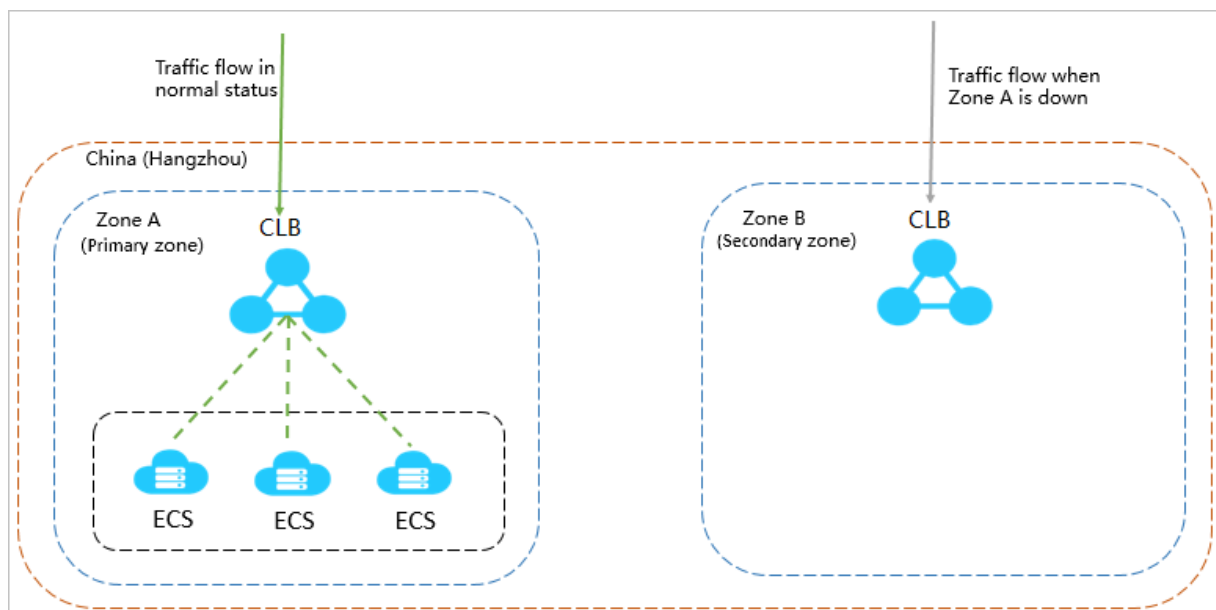
To provide more stable and reliable load balancing services, Alibaba Cloud allows you to deploy CLB instances across multiple zones in most regions for disaster recovery. Specifically, you can deploy a CLB instance in two zones within the same region. One zone is the primary zone, while the other zone is the secondary zone. If the primary zone fails or becomes unavailable, the CLB instance will fail over to the secondary zone in about 30 seconds. When the primary zone recovers, the CLB instance will automatically switch back to the primary zone.

We recommend that you create a CLB instance in a region that has multiple zones for zone-disaster recovery. We recommend that you plan the deployment of backend servers based on your business needs. In addition, we recommend that you add at least one backend server in each zone to achieve the highest load balancing efficiency.

As shown in the following figure, ECS instances in different zones are attached to a single CLB instance. In normal cases, the CLB instance distributes inbound traffic to ECS instances both in the primary zone (Zone A) and in the secondary zone (Zone B). If Zone A fails, the CLB instance distributes inbound traffic only to Zone B. This deployment mode helps avoid service interruptions caused by zone-level failure and reduce latency.




Assume that you deploy all ECS instances in the primary zone (Zone A) and no ECS instances in the secondary zone (Zone B) as shown in the following figure. If Zone A fails, your services will be interrupted because no ECS instances are available in Zone B. This deployment mode achieves low latency at the cost of high availability.



## Geo-disaster recovery

You can deploy CLB instances in different regions and attach ECS instances of different zones within the same region to a CLB instance. You can use DNS to resolve domain names to service addresses of CLB instances in different regions for global load balancing purposes. When a region becomes unavailable, you can temporarily stop DNS resolution within that region without affecting user access.



## 6.Terms

This topic introduces the terms used in CLB.

Term	Description
CLB	CLB distributes traffic across ECS instances. CLB provides Layer-4 and Layer-7 load balancing.
CLB instance	A load-balancing instance in CLB. To get started with CLB, you must create a CLB instance.
Endpoint	An IP address assigned to a CLB instance. The IP address can be either public or private, depending on the type of the CLB instance. You can resolve a domain name to a public IP address of a CLB instance to provide external services.
Listener	A listener distributes requests to backend servers. Each CLB instance must have at least one listener.
Backend server	A backend server is an ECS instance that receives client requests distributed by a CLB instance.
Default server group	<p>A group of ECS instances that process distributed requests.</p> <p>If a listener is not configured with any VServer group or primary/secondary server group, the listener distributes traffic to the backend servers in the default server group.</p>
VServer group	<p>A group of ECS instances that process distributed requests.</p> <p>You can create multiple VServer groups for different listeners of a CLB instance to specify traffic distribution with specific listeners.</p>
Primary/secondary server group	Each primary/secondary server group contains two ECS instances, where one acts as the primary server and the other acts as the secondary server. If the primary server is detected unhealthy, new requests are automatically distributed to the secondary server.