

Server Load Balancer

Product Introduction

Product Introduction

What is Server Load Balancer

Alibaba Cloud Server Load Balancer is a traffic distribution control service that distributes the incoming traffic among multiple Elastic Compute Service (ECS) instances according to the configured forwarding rules. It expands the service capabilities of the application and increases the availability of the application.

By setting a virtual service IP address, Server Load Balancer virtualizes the ECS instances located in the same region into a high-performance and highly-available application service pool. Client requests are distributed to the ECS instances in the cloud server pool according to the defined forwarding rules.

Server Load Balancer checks the health status of the ECS instances in the cloud server pool and automatically isolates any ECS instances with an abnormal status. This eliminates the single point of failure (SPOF) of an ECS instance and improves the overall service capability. Additionally, Server Load Balancer also provides the capability of defending DDoS attacks, which enhances security of the application.

Components

Server Load Balancer consists of the following components:

Server Load Balancer instances: A Server Load Balancer instance is a running load balancing service that receives and distributes the incoming traffic to the backend servers.


To use the Server Load Balancer service, you must create a Server Load Balancer instance with at least one listener and two ECS instances configured.

Listeners: A listener checks the client requests and forwards the requests to the backend servers. It also performs health check on the backend servers.

Backend servers: Backend servers are the ECS instances added to a Server Load Balancer instance to process the distributed requests. You can group the ECS instances hosting

different applications or functioning different roles into different server groups.

As shown in the following figure, after the Server Load Balancer instance receives a client request, the listener forwards the request to the corresponding backend ECS instances according to the configured listening rules.



Benefits

High availability

Server Load Balancer is designed to work in the full-redundancy mode without SPOF. Server Load Balancer supports local and cross-region disaster tolerance. When Server Load Balancer is used together with DNS, the service availability is up to 99.95%.

Scalability

You can scale your service based on the application load, without interrupting services continuity.

Low cost

Compared with the traditional hardware load balancing system, Server Load Balancer reduces the cost by 60%.

Security

Combined with Alibaba Cloud Security, Server Load Balancer can defend against up to 5 Gbps DDoS attacks, such as HTTP flood and SYN flood attacks.


Architecture

Alibaba Cloud provides the layer-4 (TCP protocol and UDP protocol) and layer-7 (HTTP protocol and HTTPS protocol) load balancing services. Deployed in clusters, Server Load Balancer can synchronize sessions to protect the ECS instances from single points of failure (SPOFs). This improves redundancy and guarantees the service stability.


Alibaba Cloud layer-4 and layer-7 load balancing services are realized in the following ways, respectively:

Layer 4 uses the open source software Linux Virtual Server (LVS) with keepalived to achieve load balancing, and also makes some customization to it according to the cloud computing requirements.


Layer 7 uses Tengine to achieve load balancing. Tengine, a Web server project based on Nginx, adds a wide range of advanced features dedicated for high-traffic websites.



As shown in the following figure, the layer-4 load balancing in each region is actually run in a cluster of multiple LVS machines. The cluster deployment model strengthens the availability, stability, and scalability of the load balancing services in abnormal circumstances.



Additionally, each LVS machine in the LVS cluster uses multicast packets to synchronize sessions to other LVS machines. As shown in the following figure, session A established on LVS1 is synchronized to other LVS machines after three packets are transferred. In normal situations, the session request is sent to LVS1 as the solid line shows. If LVS1 is abnormal or being maintained, the session request will be sent to other machines working normally, as the dotted line shows. In this way, you can perform hot upgrades, machine failure maintenance, and cluster maintenance without affecting business applications.



Note: If a connection is not established (three-way handshake is not completed), or a connection has been established but the session synchronization mechanism is not triggered, the hot upgrade does not guarantee that the connection is not interrupted and the client needs to re-initiate the connection.

Features

Alibaba Cloud Server Load Balancer has the following features:

Supported protocol

Alibaba Cloud provides both layer-4 (TCP and UDP) and layer-7 (HTTP and HTTPS) load balancing services.

Health check

Through health check on backend ECS instances, Server Load Balancer can automatically block abnormal ECS instances and distribute requests to them when they become normal.

Session persistence

Server Load Balancer supports session persistence. You can set listening rules to forward a session request from a client to the same backend ECS instance during the session lifecycle.

Scheduling algorithm

Server Load Balancer supports the following scheduling algorithms:

Round robin

Requests are distributed across the backend ECS servers sequentially.

Weighted round robin (WRR)

You can set a weight for each backend server. Servers with higher weights receive more requests than those with lower weights.

Weighted least connections (WLC)

In addition to weights set for backend ECS servers, the number of connections to servers is also considered. The server with a higher weight value will receive a larger percentage of live connections at any one time. If the weights are the same, the system directs network connections to the server with the smallest number of established connections.

Domain name/URL-based forwarding

For Layer-7 (HTTP and HTTPS) protocols, Server Load Balancer forwards traffic to different VServer groups based on domain names or URLs.

Multiple zones

To provide a more stable and reliable load balancing service, Server Load Balancer can be deployed in multiple zones in most regions. If the primary zone becomes abnormal, the backup zone automatically takes over the load balancing service from the faulty zone.

Access control

You can set a whitelist to control which IP addresses can access Server Load Balancer.

Security

Server Load Balancer supports application firewalls and HTTP flood protection. Combined with Alibaba Cloud Security, the system can defend against up to 5 Gbps DDoS attacks.

Certificate management

Server Load Balancer service provides Certificate Management for the HTTPS protocol listening. With Certificate Management, you do not need to upload certificates to backend ECS instances. Deciphering is performed on Server Load Balancer to reduce the CPU overheads of backend ECS instances.

Bandwidth control

You can set the bandwidth peak for each listener based on the service that the application can provide.

Instance type

You can choose to create an Internet or intranet Server Load Balancer service. The system will assign a public IP address or private IP address accordingly.

Monitoring

With the Monitor function, you can get the real-time status of your Server Load Balancer.

Management methods

You can manage Server Load Balancer instances through various methods, such as the Server Load Balancer console, API, and SDK.

Scenarios

Server Load Balancer is applicable to the following scenarios:

Scenario 1: Distribute traffic load for high-traffic applications

If your application traffic is high, you can use Server Load Balancer to distribute the traffic to multiple ECS instances. Additionally, you can use the session persistence feature to forward session requests from a client to the same backend ECS instance to improve access efficiency.

Scenario 2: Expand service capability for applications

You can extend service capabilities by adding and removing backend ECS instances at any time, depending on your business needs. It is applicable to Web and App applications.

Scenario 3: Eliminate single point of failure (SPOF)

With the health check feature, Server Load Balancer automatically blocks unhealthy ECS instances and distribute requests to healthy ECS instances, eliminating any single point of failure.


Scenario 4: Disaster tolerance in one region (multi-zone disaster tolerance)

To provide more reliable services, multiple zones for Server Load Balancer are deployed in most regions.

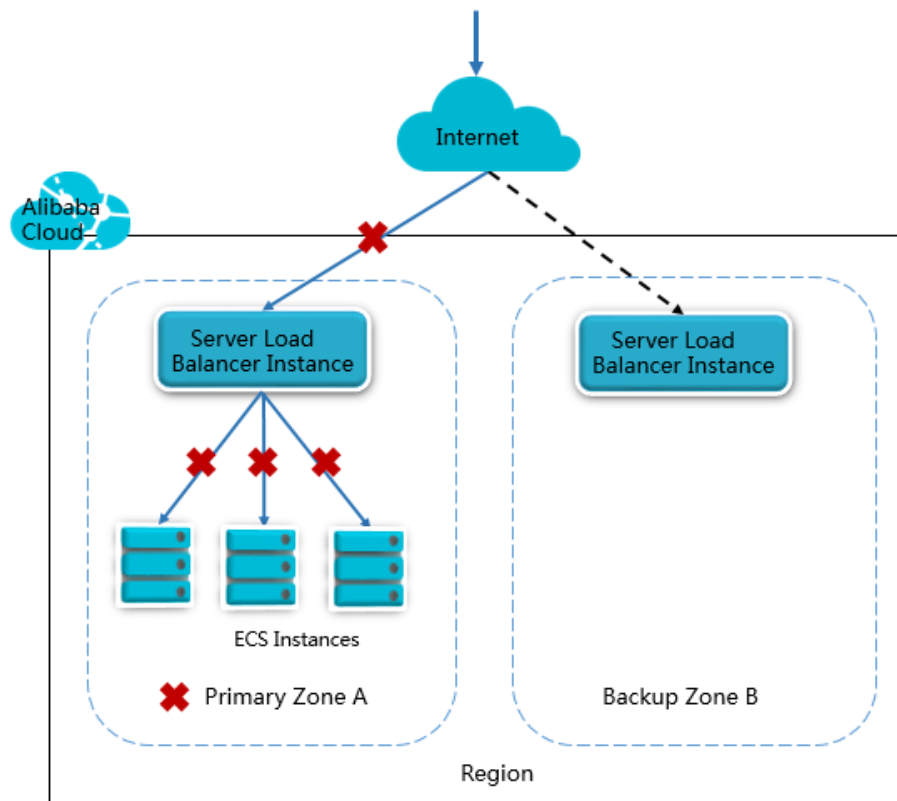
If a primary zone becomes unavailable, Server Load Balancer rapidly switches to a backup zone to restore its service capabilities within 30 seconds. When the primary zone becomes available, Server Load Balancer automatically switches back to the primary zone.

We recommend that you create a Server Load Balancer instance in a region with multiple zones for disaster tolerance. You can deploy ECS instances as needed. It is a best practice to add at least one ECS instance in each zone, which can achieve load balancing featuring higher availability and lower latency.

As shown in the following figure, ECS instances in different zones are added to the Service Load Balancer instance. When primary zone A works normally, traffic is distributed to ECS instances in the primary zone A, as the blue line shows. When primary zone A becomes unavailable, traffic is distributed to ECS instances in the backup zone, as the black dotted line shows. This avoids service interruption because of failure of a single zone, and also reduces latency because you can freely choose zones with lower latency.

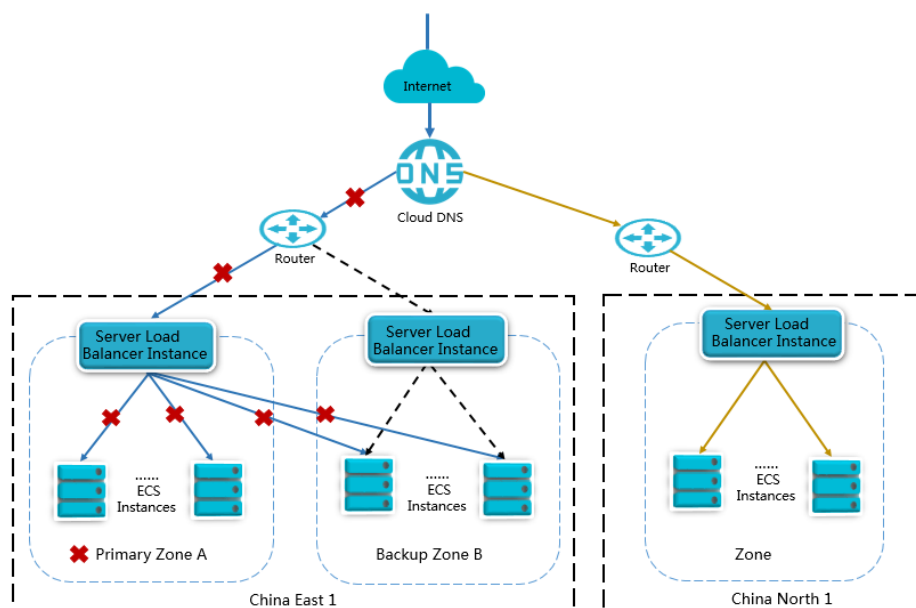


If you add all the ECS instances to the primary zone and have no ECS instances in the backup zone, as shown in the following figure, low latency is achieved at the expense of high availability. In this situation, when the primary zone is unavailable, Server Load Balancer switches to the backup zone. However, there are no ECS instances to handle the distributed requests in the backup zone.



Scenario 5: Cross-region disaster tolerance

Combined with DNS, you can achieve cross-region disaster tolerance by deploying Server Load Balancer in different regions and using DNS to resolve the domain name to the IP addresses of the Server Load Balancer instances. When a region becomes unavailable, you need to stop domain name resolution for the region so that user access is not affected.



Limits

| Resource | Limit | Exception |
|---------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------|
| ECS instances added to a Server Load Balancer instance | Must have at least one ECS instance, and the region must be the same for the ECS instance and Server Load Balancer instance. | No |
| Types of a Server Load Balancer instance | Internet/intranet | No |
| Billing methods | By traffic or bandwidth | No |
| Public network bandwidth range for a Server Load Balancer instance (Pay by bandwidth) | 1 to 5000 MB, the default is 1 MB | No |
| Default quota of Pay-As-You-Go instances | 60 | Submit a ticket to apply for more |
| Server Load Balancer instance name | The name can be 1-80 characters in length, and include letters, digits, hyphens (-), backslashes (/), periods (.), and underscores (_). | No |
| Number of listeners added in a Server Load Balancer instance | 50 | No |
| Maximum number of domain name/URL forwarding rules added to a layer-7 listener | 20 | Submit a ticket to apply for more |
| Supported protocols | HTTP/HTTPS/TCP/UDP | No |
| Front-end/backend port range used in a listener(public cloud) | 1 to 65535 | No |
| Front-end/backend port range used in a listener (Financial Cloud) | 80, 443, 2800 to 3300, 5000 to 10000, 13000 to 14000 | No |
| Supported scheduling algorithms | Round robin, weighted round robin (WRR), weighted least connections (WLC) | No |
| Cookie processing in HTTP | insert/server, the default is | No |

| | | |
|----------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| session persistence | insert | |
| Cookie timeout in HTTP session persistence | 1 to 86400, the default is 3600 | No |
| Cookie name in HTTP session persistence | <p>The name can be 1 to 200 characters in length, and must comply with RFC2965.</p> <p>It can only contain ASCII English letters and digits.</p> <p>It cannot contain commas, semicolons, spaces, or begin with a dollar symbol (\$).</p> | No |
| Health check port for HTTP protocol | 1 to 65,535 (the default is the backend server port) | No |
| Health check domain name for HTTP protocol | The length can be 1 to 80 characters in length, and include letters, digits, hyphens (-), and periods (.). | No |
| Health check URI for HTTP protocol | The URU can be 1-80 characters in length, and include letters, digits, hyphens (-), backslashes (/), periods (.), percent signs (%), question marks (?), number signs (#), and ampersands (&). | No |
| Health check timeout for HTTP protocol | 1 to 300s, the default is 5 | No |
| Health check interval for HTTP protocol | 1 to 50s, the default is 2 | No |
| Healthy threshold for HTTP protocol health check | 2 to 10, the default is 3 | No |
| Unhealthy threshold for HTTP protocol health check | 2 to 10, the default is 3 | No |
| Session timeout for TCP protocol | 1 to 3600s | No |
| Health check port for TCP protocol | 1 to 65,535 (the default is the backend server port) | No |
| Health check timeout for TCP protocol | 1 to 300s, the default is 5 | No |
| Health check interval for HTTP protocol | 1 to 50s, the default is 2 | No |
| Healthy threshold for TCP protocol health check | 2 to 10, the default is 3 | No |
| Unhealthy threshold for TCP protocol health check | 2 to 10, the default is 3 | No |

| | | |
|-------------------------------------------------------------------------|------------------------------------------------------|--------------------------------------------------------------------------------------|
| Health check port for UDP protocol | 1 to 65,535 (the default is the backend server port) | No |
| Health check timeout for UDP protocol | 1 to 300s, the default is 10 | No |
| Health check interval for UDP protocol | 1 to 50s, the default is 10 | No |
| Healthy threshold for TCP protocol health check | 2 to 10, the default is 3 | No |
| Unhealthy threshold for UDP protocol health check | 2 to 10, the default is 3 | No |
| Number of backend ECS instances that can be added or deleted in batches | 20 | No |
| API access frequency for a single key | 5,000 times/day | Currently, an automatic process is not available. Contact customer manager for help. |
| Maximum number of certificates that can be uploaded by a single user | 100 | Currently, an automatic process is not available. Contact customer manager for help. |

Terms

| Term | Description |
|-------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Server Load Balancer | Alibaba Cloud Server Load Balancer is a traffic distribution control service. It distributes incoming application traffic among multiple ECS instances according to the configured scheduling algorithm and listening rules. |
| Server Load Balancer Instance | A Server Load Balancer instance is a running instance of the Server Load Balancer service. To use Server Load Balancer, you must first create a Server Load Balancer instance. The instance ID is a unique identifier for the Server Load Balancer instance. |
| Server Load Balancer IP | The IP address allocated to the Server Load Balancer service after a Server Load Balancer instance is created. According to the instance type, the IP address is either a public IP or a |

| | |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | <p>private IP. You can resolve a domain name to a public IP address to provide external services.</p> |
| Listener | <p>A listener defines how the incoming requests are distributed. You must add at least one listener to a Server Load Balancer instance.</p> |
| Backend Server | <p>The ECS instances that process the distributed requests.</p> |
| VServer Group | <p>A group of ECS instances that process the distributed requests. Different listeners can use different VServer groups, which allows you to maintain the request distribution in the listener dimension.</p> |
| Multiple Zones | <p>Multiple zones for Server Load Balancer have already been deployed in most regions for better disaster tolerance.</p> <p>By default, the instance in the primary zone is used to provide load balancing service, when the instance in the primary zone is unavailable, the instance automatically switches to the backup zone to continue providing service and then switches back to the primary zone when it becomes available. This increases local availability.</p> |