

Server Load Balancer

Product Introduction

Product Introduction

What is Server Load Balancer ?

Alibaba Cloud Server Load Balancer is a traffic distribution control service, which distributes the incoming application traffic among multiple ECS instances according to the scheduling algorithm and listening rules.

By setting a virtual IP address, the Server Load Balancer service virtualizes the ECS instances located in the same region into a high-performance, and highly available application service pool. The client requests are distributed to the cloud server pool according to the defined listening rules. This increases the fault tolerance of your applications.


Besides, the Server Load Balancer service checks the health status of the ECS instances in the cloud server pool and automatically isolates the ECS server of the abnormal state. This resolves the single point of failure (SPOF) problem and improves the overall service capability of the application.

In addition to the standard load balancing capabilities, the TCP and HTTP listening also can defend the DDoS attacks, which enhances the protection capability of the application servers.

Components of Server Load Balancer

The Server Load Balancer service consists of three parts: a Server Load Balancer instance, one or more listeners and multiple backend servers.

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.



Server Load balancer Instance: If you want to use the Server Load Balancer service, you must create a Server Load Balancer instance. You can add multiple listeners and backend servers to a Server Load Balancer instance.

Listener: Before using the Server Load Balancer service, you must add at least one listener to the Server Load Balancer instance, which defines how the client requests are forwarded to the backend servers.

Backend server: The ECS instances added to the Server Load Balancer instances, which are the backend servers used to process the distributed requests. You can add the ECS instances separately to the backend server pool, or add them in batch through a VServer group or master-slave server group.

By default, the backend servers are maintained in the Server Load Balancer instance dimension. All listeners can only forward requests to the same ECS instances with the same backend port configured in the listeners. With the VServer group function, you are allowed to maintain the backend servers in the listener dimension. You can create different VServer groups for different listeners, that is, the listeners in a Server Load Balancer instance can forward requests to different backend servers with different ports.

In addition, the Layer-7 load balancing service supports configuring domain name or URL forwarding rules. The listener can forward the requests from different domain names or URLs to different VServer groups.

Benefits

High availability

Designed to work in full-redundancy mode without single points of failure (SPOF). Server Load Balancer supports local and cross-region disaster tolerance when used together with DNS, delivering service availability of up to 99.95%.

Server Load Balancer can scale its service flexibly based on the application load without interrupting external services during traffic fluctuation.

Low cost

Server Load Balancer is 60% more cost-efficient than traditional hardware load-balancing systems. By providing free access to private network instances without generating any O&M cost, the service completely removes the need to purchase expensive load-balancing equipment.


Security

Combined with Alibaba Cloud Security, Server Load Balancer can defend against DDoS attacks, such as CC and SYN flood attacks.


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 the redundancy and guarantees the service stability.

Layer 4 uses the open source software LVS (Linux Virtual Server) 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 is a Web server project sponsored by Taobao. Based on Nginx, Tengine 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 consisting of multiple LVS machines. The cluster deployment model greatly guarantees the availability, stability and scalability of the load balancing services in abnormal circumstances.



Additionally, the LVS machine in the LVS cluster uses multicast packets to synchronize sessions to other LVS machines. As shown in the following figure, the 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 becomes abnormal or is being maintained, the session request will be sent to other machines that work 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: When the 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.

Supported protocol

The 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 again when they become normal.

Session persistence

Server Load Balancer supports session persistence. You can set listening rules to forward the 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 group of 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 less weights.

Weighted least connections (WLC)

In addition to the weight set to each backend ECS server, the number of connections to the client is also considered. The servers 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 least 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 the domain names or URLs.

Multiple zones

To provide a more stable and reliable load balancing service, Alibaba Cloud Server Load Balancer has deployed multiple zones in most regions. When the primary zone is 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 the Server Load Balancer service.

Security

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

Certificate management

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

Balancer in order 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 a public or a private Server Load Balancer service. The system will assign a public IP address or a private IP address accordingly.

Monitoring

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

Management methods

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

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 the session requests from a client to the same backend ECS instance to improve the access efficiency.

Scenario 2: expand service capability for applications

You can extend the 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 the single point of failure (SPOF)

With the health check feature, the Server Load Balancer service will automatically block the unhealthy ECS instances and distribute requests to the healthy ECS instances to eliminate the single point of failure.


Scenario 4: disaster tolerance in one region

To provide more reliable services, Alibaba Cloud Server Load Balancer has already deployed multiple zones in most regions.

When the primary zone is unavailable, the Server Load Balancer service rapidly switches to the backup zone to restore its service capabilities within 30s. When the primary zone becomes available, the Server Load Balancer service will automatically switch back to the primary zone.


We recommend that you create a Server Load Balancer instance in the region with multiple zones for disaster tolerance. Additionally, you can get a higher-availability or lower-latency load balancing service by reasonably deploying ECS instances. It is a best practice that adding at least one ECS instance in each zone.

As shown in the following figure, when the primary zone A works normally, the traffic is distributed to the ECS instances in the primary zone A as the blue line shows. When the primary zone A becomes unavailable, the traffic is distributed to the ECS instances in the backup zone as the black dotted line shows. This avoids the business interruption due to the failure of a single zone, and also reduces the latency between the products in different zones.




If you add all the ECS instances to the primary zone and no ECS instances in the backup

zone, as shown in the following figure, the low latency is achieved at the expense of high availability. In this situation, when the primary zone is unavailable, the Server Load Balancer will switch to the backup zone. However, there is no ECS instances to handle the distributed requests.



Scenario 5: cross-region disaster tolerance

Combined with DNS, you can achieve cross-region disaster tolerance by deploying the Server Load Balancer services in different regions and using DNS to resolve the domain to the IP addresses of the Server Load Balancer instances. When a region becomes unavailable, you just need to stop domain name resolution without affecting user access.



Item	Limitation	Exception
ECS instances added to a Server Load Balancer instance	Must have one ECS instance at least, and the region of the ECS instance and the Server Load Balancer instance must be the same.	No
Types of a Server Load Balancer instance	Public/Private	No
Billing methods	By traffic or by bandwidth	No
Public bandwidth range for a Server Load Balancer instance (Pay by bandwidth)	1-5000 Mbps , default is 1 Mbps	No
Peak public bandwidth range for a single listener (Pay by bandwidth)	1-5000 Mbps Note: If the billing method is PayByTraffic, no restrictions on the bandwidth for the listeners.	No
Default quota of Pay-As-You-Go instances	60	Submit a ticket to apply for more
Restrictions on Server Load Balancer instance name	The length is 1-80 characters, including letters, digits, hyphen (-), backslash (/) , period (.) and underscore (_)	No
The number of listeners added in a Server Load	50	No

Balancer instance		
The number of the domain/URL forwarding rules added to a layer-7 listener	20	Yes
Supported protocols	HTTP/HTTPS/TCP/UDP	No
Front-end/backend port range used in a listener	1-65535	No
Front-end/backend port range used in a listener (Financial Cloud)	80 , 443 , 2800-3300 , 5000-10000 , 13000-14000	No
Supported scheduling algorithms	Round robin, Weighted round robin (WRR), Weighted least connections (WLC)	No
Cookie processing in HTTP session persistence	insert/server , default is insert	No
Cookie timeout for HTTP protocol	1-86400, default is 3600	No
Cookie name restrictions	The length is 1-200 , and must comply with RFC2965. This means that it only contain ASCII English letters and digits, and cannot contain commas, semicolons, spaces, or begin with a dollar symbol "\$"	No
Health check port for HTTP protocol	1 – 65,535 (default is the backend server port)	No
Health check domain for HTTP protocol	The length is 1-80 characters, including letters, digits, hyphen (-), period (.)	No
Health check URI for HTTP protocol	The length is 1-80 characters, including letters, digits, ' - ' , ' / ' , ' : ' , ' % ' , ' ? ' , ' # ' , ' & '	No
Health check timeout for HTTP protocol	1-300s, default is 5	No
Health check interval for HTTP protocol	1-50s, default is 2	No
Healthy threshold for HTTP protocol health check	2-10, default is 3	No
Unhealthy threshold for HTTP protocol health check	2-10, default is 3	No
Session timeout for TCP protocol	1-3600s	No

Health check port for TCP protocol	1 – 65,535 (default is the backend server port)	No
Health check timeout for TCP protocol	1-300s, default is 5	No
Health check interval for HTTP protocol	1-50s, default is 2	No
Healthy threshold for TCP protocol health check	2-10, default is 3	No
Unhealthy threshold for TCP protocol health check	2-10, default is 3	No
Health check port for UDP protocol	1 – 65,535 (default is the backend server port)	No
Health check timeout for UDP protocol	1-300s, default is 10	No
Health check interval for UDP protocol	1-50s, default is 10	No
Healthy threshold for TCP protocol health check	2-10, default is 3	No
Unhealthy threshold for UDP protocol health check	2-10, default is 3	No
Number of backend ECS instances that can be added or deleted in batch	20	No
API access frequency limit for a single key	5,000 times/day	No automatic process is available now. Contact customer service or Business Development for help.
Maximum number of certificates uploaded by a single user	100	No automatic process is available now. Contact customer service or Business Development for help.

Term	Description
Server Load Balancer	Alibaba Cloud Server Load Balancer is a traffic distribution control service, which distributes the 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 the unique identifier of the

	Server Load Balancer instance.
Server Load Balancer IP	The IP address allocated to the Server Load Balancer service after creating a Server Load Balancer instance. According to the instance type, the IP address is either a public IP or a private IP. You can resolve a domain name to the public IP address to provide external services.
Listener	A listener defines how the incoming requests are distributed. You must add at least one listener to a Server Load Balancer instance.
Backend Server	The ECS instances that process the distributed requests.
VServer Group	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.
Multiple Zones	Server Load Balancer has already deployed multiple zones in most regions for better disaster tolerance. By default, the instance in the primary zone is used to provide the load balancing service. However, when the primary zone is unavailable, the instance will automatically switch to the backup zone to continue providing service and switch back to the primary zone when the primary zone become available. This greatly increases the local availability.

Date	Event
February 11, 2014	Server Load Balancer is commercially available and is billable in PayByTraffic mode.
March 10, 2014	An updated version of the Server Load Balancer console is released with the following changes: correction of basic instance information, restoration of the Create Instance button on the instance list page, correction of the Help Center link, and addition of tips to the metrics page.
April 25, 2014	Server Load Balancer is released in North China 2 (Beijing) region.
May 8, 2015	Server Load Balancer is released in Hong Kong region.
May 19, 2014	The Server Load Balancer API is released. The API gives more freedom and ease-of-use for activating and managing Server Load Balancer

	and creating third-party applications based on the computing resources of Alibaba Cloud.
August 29, 2014	Server Load Balancer is released in South China 1 (Shenzhen) region.
November 14, 2014	The Alibaba Cloud Server Load Balancer system is upgraded. A local high-availability solution capable of local disaster tolerance is deployed first in East China 1 (Hangzhou) and North China 1 (Qingdao) regions.
January 29, 2015	The Server Load Balancer HTTPS protocol is supported to transmit encrypted data securely.
February 3, 2015	The default maximum number of Server Load Balancer instances that a user can purchase is increased from 10 to 30, and the maximum number of listeners of a single instance is increased from 10 to 50 so that users can deploy more instances to meet their growing requirements.
March 24, 2015	The Server Load Balancer console supports configuring performance monitoring alarms so that users can quickly understand how to use the performance indicators of Server Load Balancer.
April 24, 2015	VPC-capable Server Load Balancer is released.
July 1, 2015	UDP-capable Server Load Balancer is released, and Layer-4 TCP monitoring supports HTTP health check.
December 18, 2015	Multi-zone Server Load Balancer is released in order to support disaster tolerance across machine rooms in the same region.
January 27, 2016	The security protection feature of Server Load Balancer is released. The feature involves application-layer web application firewall (WAF) protection and Challenge Collapsar (CC) protection.
March 5, 2016	Server Load Balancer supports domain names and URL forwarding.
June 15, 2016	The product documentation of Server Load Balancer is revised.