

An overview of the different load balancing solutions that are available on Google Cloud.

(/load-balancing/docs/load-balancing-overview)

A guide to help you determine which Google Cloud load balancer best meets your needs.

(/load-balancing/docs/choosing-load-balancer)

A guide to selecting the appropriate load balancer based on your application needs.

(/load-balancing/docs/features)

An overview of Internal HTTP(S) Load Balancing, including traffic types, scheme, and scope, as well as use cases and architecture.

(/load-balancing/docs/l7-internal/index)

Internal HTTP(S) Load Balancing supports advanced traffic management functionality that enables you to use traffic steering, traffic actions, and traffic policies.

(/load-balancing/docs/l7-internal/traffic-management)

An introduction to the information that you need to configure external HTTP(S) Load Balancing.  
(</load-balancing/docs/https/index>)

An overview of Internal TCP/UDP Load Balancing, which enables you to run and scale your services behind a private load balancing IP address that is accessible only to your internal virtual machine (VM) instances.  
(</load-balancing/docs/internal/index>)

Provides scenarios for accessing an internal TCP/UDP load balancer in your Virtual Private Cloud (VPC) network from a connected network.  
(</load-balancing/docs/internal/internal-lb-and-other-networks>)

Provides one method of increasing availability, while also giving you greater control over how to manage your workload when your primary backend VMs aren't healthy.  
(</load-balancing/docs/internal/failover-overview>)

You can use an internal TCP/UDP load balancer as the next gateway to which packets are forwarded along the path to their final destination by setting the load balancer as the next hop in a custom static route.

(</load-balancing/docs/internal/ilb-next-hop-overview>)

An overview of Network Load Balancing, including protocols, scheme, and scope, as well as architecture.

(</load-balancing/docs/network/index>)

SSL Proxy Load Balancing terminates user SSL (TLS) connections at the load balancing layer, and then balances the connections across your instances by using the SSL or TCP protocols. The SSL proxy is intended for non-HTTP(S) traffic.

(</load-balancing/docs/ssl/index>)

TCP Proxy Load Balancing allows you to use a single IP address for all users around the world. The TCP proxy load balancer automatically routes traffic to the instances that are closest to the user. The TCP proxy is intended for non-HTTP(S) traffic.

(</load-balancing/docs/tcp/index>)

An overview of backend services for your load balancer. A backend service contains configuration values for load balancing services.

(/load-balancing/docs/backend-service)

A forwarding rule and its corresponding IP address represent the frontend configuration of a Google Cloud load balancer.

(/load-balancing/docs/forwarding-rule-concepts)

An overview of health checking mechanisms specific to Google Cloud and its load balancers.

(/load-balancing/docs/health-check-concepts)

IPv6 termination enables you to handle IPv6 requests from your users and proxy them over IPv4 to your backends.

(/load-balancing/docs/ipv6)

An introduction to using network endpoint groups (NEGs) as the backend for a backend service. The primary use case for this configuration is deploying containers on your VMs so that you can run services in the containers.

(/load-balancing/docs/negs/index)

SSL policies give you the ability to control the features of SSL that your SSL proxy load balancer or external HTTP(S) load balancer negotiates with clients.

(/load-balancing/docs/ssl-policies-concepts)

HTTP(S)-based load balancers use a URL map to direct incoming requests to backend services and backend buckets.

(/load-balancing/docs/url-map-concepts)