

CumuloLogic Load Balancing Service

Load balancers play a critical role in scaling applications in the cloud. Load balancers distribute the incoming traffic to several web server or application server nodes or clusters and provide a mechanism to scale your applications based on the workload. Unlike on-premise load balancers, cloud applications use software load balancers.

CumuloLogic platform provides a specialized, low cost and scalable load balancer for running on any Infrastructure-as-a-Service (IaaS) cloud. CumuloLogic load balancing service offers all the features of an on-premise load balancer, in addition to providing benefits, such as easy deployment, management, automatic node discovery for managed applications, etc. It enables you to achieve even greater fault tolerance in your applications, seamlessly providing the amount of load balancing capacity needed in response to incoming application traffic. CumuloLogic load balancing service detects unhealthy nodes and automatically reroutes traffic to healthy nodes until the unhealthy nodes have been recovered.

CumuloLogic load balancing service is extremely simple to deploy and it provides high availability, resiliency, and performance optimization functions to meet the requirements of mission critical applications.

With CumuloLogic load balancing service, you can build and deliver a fully managed and highly available "Load Balancer-as-a-Service," and design application architectures to fit specific availability, scalability and performance needs.

CumuloLogic load balancing service is integrated with CumuloLogic Platform for developers to deploy any language application and scale it using the load balancing service.

CumuloLogic Load Balancing Service At-A-Glance

Features supported by CumuloLogic load balancing service vary depending on the load balancer engine selected. Below are some common features:

Provisioning – Single-click or single-API call to provision the desired load balancer instance: size with configuration and performance parameters, and rules for auto update and patches.

Monitoring – Visibility into key functional and operational metrics of the load balancer engine and instance, including CPU, memory, I/O, network traffic, and response time.

Recovery: Automated recovery of load balancer from a service or instance failure. In the event of a service failure, the health agent will recover the service seamlessly, and it will re-provision a fully configured load balancer instance to replace the failed instance.

Automated Software Patching – Optional feature to apply minor updates and patches to the load balancer engine. Users can choose when the updates are to be performed based on their application maintenance downtime.

Optimization – Set of default performance parameters for the chosen load balancer engine and the instance size. Users can modify the performance parameters to suit the application

workloads in real-time and in most cases without shutting down the running load balancer instances or rebooting.

Service Benefits:

Easy-to-use with single-click deployment: New load balancer instances can be quickly provisioned either using the user interface or a single API call. The load balancer service can be launched separately and then bound to a running application or can be orchestrated to launch during application deployment.

Low TCO: Load balancer instances are fully managed and monitored, eliminating most manual tasks and substantially lowering the cost of application operations.

Security at multiple levels: Load balancer instances are secured using the firewall settings and security groups of the chosen IaaS cloud, allowing users to control remote access to all load balancers instances. Additionally, load balancer instances can be configured to use secured connections only.

Any-cloud: CumuLogic's platform abstracts the underlying APIs of the IaaS cloud or vSphere environment, allowing users to deploy the platform and the load balancing service on any supported private or public cloud.

Autoscaling: For CumuLogic platform managed-applications, CumuLogic load balancing will automatically adjust the Load balancing configurations to match the autoscale up/down events in the application cluster.

Managing Node failures: CumuLogic load balancer can detect unhealthy load-balanced instances, and it will mark the node and automatically spread the traffic across the remaining healthy nodes.

Improved fault-tolerance for your application: You get improved fault tolerance by having CumuLogic load Balancer support your applications, as the load balancer can automatically distribute traffic across multiple nodes and multiple datacenters, and ensure that only healthy nodes receive traffic.

Management Access: CumuLogic load Balancer provides full management access via a UI console as well as RESTful APIs.

Architecture diagram

