1.5 Load Balancing

Micron21 - 2025-03-18 - Section 1: Networks & Security

Distributed Layer 4 Load Balancer

Within our mCloud portal users are provided by default access to our Distributed Layer 4 Load Balancer LBaaS to load balance layer 4 TCP and UDP traffic. Our layer 4 load balancer within mCloud Operates at the Transport Layer (Layer 4), making it protocol-agnostic to application-layer details, focusing primarily on TCP and UDP traffic. Incorporating our layer 4 mCloud load balancer will enhance the scalability and reliability by distributing client requests to multiple backend servers using load balanced algorithms which will provide fault tolerance to your virtual machine instances by rerouting traffic when instances fail, this Improves performance through efficient resource utilisation and network traffic management. Unlike centralised load balancers, our mCloud distributed model places load-balancing capabilities closer to the client and backend resources, minimising latency and enhancing throughput.

Layer 7 Advance Enterprise Load Balancing

For those customers who require Advance Layer 7 Load Balancing functionality for Server Load Balancing or Global Server Load Balancing our mCloud platform allows you to create dedicated licensed A10 virtual load balancers, designed to optimise application performance, reliability, and scalability, incorporating advanced features at higher OSI layers, enhanced performance optimizations, and global-level traffic management. Licences are provided at an additional cost based on total appliance throughput calculated in mbits.

The A10 SLB Enterprise Load balancer supports Layer 7 application aware routing, enabling decisions based on HTTPS headers, cookies, and application content, providing advanced health checks to verify service availability at the application level, not just network reachability. The appliance fully supports integrated SSL offloading freeing up server resources and improving performance along with supporting advanced encryption methods for secure connections including TLS 1.3

Our advanced appliance supports connection persistence which ensures session stickiness for consistent user experiences using mechanisms like source IP, cookies, or headers and Application Acceleration Features like caching,

compression, and TCP optimization enhance application delivery and reduce latency. If High Availability (HA) if required then the A10 appliance can be deployed in Active-active or active-standby configurations providing resilience, ensuring no single point of failure.

If you require Global Traffic Distribution then the GSLB function of a dedicated A10 appliance distributes traffic across multiple geographic locations, improving global application performance and availability. By providing DNS-Based Load Balancing the appliance will direct users to the best-performing or closest data center based on metrics like latency, geography, or server health, which can also be used for Disaster Recovery where traffic is automatically routes to backup locations in case of primary data center failures.