



Highly available storage for mission-critical applications



HIGH AVAILABILITY



EXTREME WORKLOADS



HIGH THROUGHPUT

Scality, the team behind the world's fastest object store and Loadbalancer.org, the engineers' choice for smarter load balancing, are partnering to bring IT teams a complete storage solution delivering high availability, performance and scalability.

Ideal workloads include AI, genomics, research, service providers, and media — in short, any critical workload with high throughput and zero downtime requirements.

Scality teams are able to resell Loadbalancer.org software, meaning the load balancer is considered at all points of the cycle — from solution design and sizing through to quoting. This creates a simple, frictionless process for partners and prospects while delivering a more complete solution.

Benefits of load balancing for object storage

Object storage environments gain resilience, scalability, and control when paired with load balancing. Together, Scality and Loadbalancer.org ensure always-on performance for the most demanding workloads.

High availability


Keeps the object store reachable even if some nodes fail or are being upgraded.

- Health checks detect unhealthy S3 endpoints and automatically remove them, protecting against hardware failure.
- Zero-downtime maintenance through rolling upgrades — connector and storage nodes can be patched without breaking application access.
- Failover handling for AZ/DC outages by rerouting requests to healthy regions.
- Protects clients from topology changes by presenting a single stable endpoint (e.g. s3.example.com).
- In multi-site scenarios, Global Server Load Balancing (GSLB) steers traffic to the nearest or still-available site.

Performance

Distributes load across many front-end nodes to handle bursty or parallel traffic.

- Enables parallelism by spreading thousands of threads evenly across connector nodes; ideal for backup or sequential operations.
- Delivers consistent throughput, avoiding “hot spots” by balancing concurrent large PUT/GET requests.
- Supports high concurrency for AI training, log ingestion, or CDN origin workloads.
- Flexible algorithms (round-robin, least-connections, or consistent hashing) optimize fairness, balance, or cache locality.



Security & observability

Provides a centralized control point before requests hit storage.

- Terminates TLS at the load balancer, offloading crypto workload and centralizing certificate management away from connector nodes.
- Enforces rate limiting and QoS per tenant or IP to prevent noisy neighbor issues.
- Consolidates logs, metrics, and traces in one layer for compliance reporting.

Product overview

Layer 4 / Layer 7 load balancing

Intelligent traffic distribution at both Layer 4 and Layer 7, ensuring responsiveness and optimal performance.

Advanced health checks

Proactively monitor and manage server health to ensure continuous availability and reliability of applications.

Intuitive web interface

Streamlined WebUI simplifies configuration and management of critical infrastructure.

Scripting

Full command line access and API support.

Automation

Remote deployment and control via scripting (API / CLI / JSON / SOAP / PowerShell), Ansible playbooks and more.

Global Server Load Balancing (GSLB)

Seamlessly distributes traffic across multiple data centers or geographically dispersed servers, optimizing delivery and supporting RING in two- or three-site configs.

Web Application Firewall (WAF)

Fortify applications against cyber threats with a robust Web Application Firewall, offering comprehensive protection.

Performance

Up to 100 Gb/s throughput per appliance with full SSL/TLS offload and termination — removing cryptographic overhead from connector nodes.

Monitoring


Onboard reporting plus tight integration with third-party tools such as Prometheus, Grafana and others.

Centralized management

Consolidated management, automation and observability across the entire ADC infrastructure via the Loadbalancer ADC Portal.

Simple sizing

Loadbalancer.org is available through Scalify in three simple sizing options: 10Gb/s, 50Gb/s and 100Gb/s. All options are delivered as software images, making it easy to run them on the same hardware vendor as the RING or ARTESCA components. Each option is provided as a high-availability pair, building HA into the solution from the very start.



Note: Professional Services are scoped as a flat fee per HA pair. An annual subscription SKU includes support and all functionality, with GSLB and WAF included at no extra cost.

For configurations above 100Gb/s or non-HA options, custom quotes are available upon request. Required hardware specs provided during design.

Learn more about the full capabilities of Loadbalancer.org

 **LOADBALANCER**

ABOUT SCALITY Scalify builds infinitely scalable, cyber-resilient S3 object storage for the world's most discerning and data-driven companies. Our patented MultiScale Architecture scales independently across all critical dimensions to meet the unpredictable demands of modern workloads — from AI/ML to backup and cloud infrastructure. Powered by CORE5 technology, Scalify delivers end-to-end cyber resilience for unbreakable protection against the threats of today and tomorrow.



SCALITY

storage software
FOR AI, BACKUP AND
CLOUD INFRASTRUCTURE

Follow us on Twitter and LinkedIn. Visit www.scality.com, or subscribe to our blog.
San Francisco • Paris • Washington, D.C • Tokyo • London