



ACCELERON LABS

Object Storage-100
(OS-100)

DATA SHEET

Acceleron OS-100 is an enterprise class software-defined object storage which can be deployed in your data center. It empowers enterprises and service providers with resilient, highly- available, petabyte-scale data storage for web, video, archives, backups, analytics and custom applications. Acceleron **OS-100** deploys on any standard x64 servers of your choice, taking immediate advantage of server and media innovation over time. Acting as a single, distributed system, the Acceleron **OS-100** scales linearly across thousands of servers, multiple sites, and an unlimited number of objects. Acceleron **OS-100** has no-single point of failure, has built-in redundancy for nodes, network interfaces, etc. and has highly resilient built-in features. Data becomes unbreakable through replication, erasure coding, immutable object-locking, and multi- geo distribution to achieve 14x9s of durability and 100% availability. Acceleron OS-100 is an object storage with integrated support for native file, object, and AWS S3 & IAM APIs interfaces — The Acceleron **OS-100** provides high performance across a variety of workloads at up to 90% lower TCO than legacy storage.

Acceleron OS-100 Use Cases

OS-100 enables both data-rich custom and packaged applications aggregating multiple workloads into a single storage environment. This eliminates silos, drives increased utilization and creates cloud-like economies of scale. **OS-100** supports use cases across many industries, including service providers, media & entertainment, public sector, financial services, healthcare, and R&D-intensive industries like manufacturing and research.

BACKUP TARGET

- Unified repository for capacity and performance
- BIG DATA ANALYTICS
- HPC / AI / ML archives
- Data lakes
- Forensics / digital evidence archives
- Genomics archives

File and Object Performance

OS-100 ensures high performance and throughput while eliminating silos. The platform's access and storage layers can scale independently from as few as three to thousands of servers, all of which can be accessed directly and concurrently. **OS-100** connectors scale performance nearly linearly to achieve tens of gigabytes per second

throughput and beyond.

Scale-Out File and Object Storage

Access all of your data at petabyte scale. **OS-100** is the only storage platform that natively supports file and object access in one consolidated system. Each access point can be independently configured for high availability

Local and Multi-Site Data Protection and

Self-Healing

OS-100 is designed to maintain 100% data availability and require no downtime during planned and unplanned events, including hardware failures, capacity upgrades and software upgrades. With **OS-100**, IT organizations have fine-grained data protection control on a per object basis, and customizable failure domains to protect against disk, server, rack, and site failures.

OS-100

Agile Deployment

With **OS-100**, you're free to deploy in ways co-located with customer application servers, on virtual machines, or on the storage servers. Deploy an **OS-100** environment on a single site or stretched over multiple sites. Choose to replicate locally, across sites, or both. **OS-100's** unified installer with silent (unattended) mode enables customers to deploy **OS-100** easily in under an hour on any of the reference platform architectures, including new 3-server entry point configurations. **OS-100** can scale up by adding new nodes to exiting cluster.

All-Flash Support

OS-100 supports all-flash storage servers, paving the way for a new level of file and object storage performance, with both higher throughput and lower latency.

that best suit your infrastructure. **OS-100** connectors and storage servers are independently scalable. Run the connectors



Always on



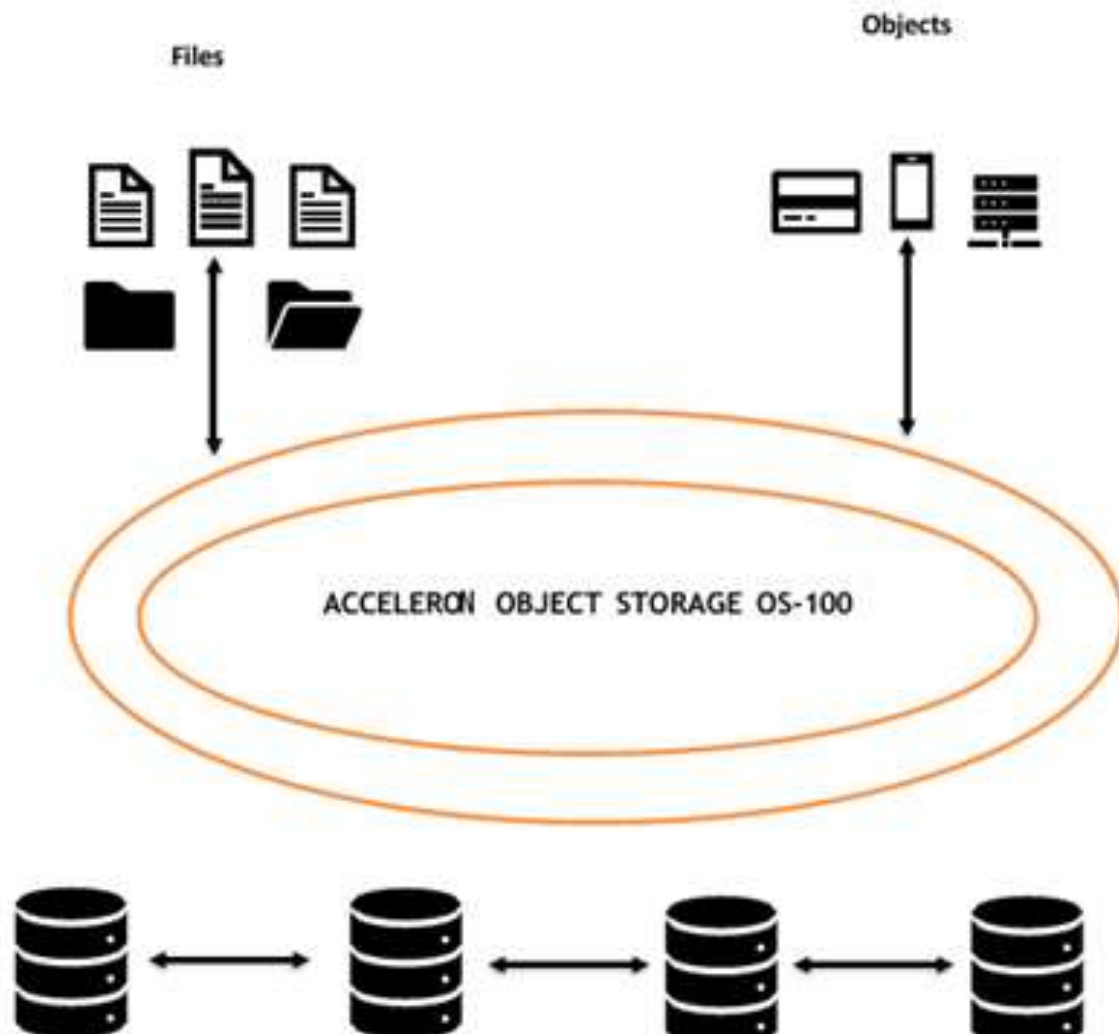
Petabyte scale



Immutable Data



Cloud Economics



Data Security

- **Data at Rest Encryption:** The object storage solution uses AES-256 encryption, a robust industry-standard algorithm, to secure data at rest, ensuring data is protected against unauthorized access while stored.
- **Data in Transit Encryption:** The solution also supports AES-256-SHA encryption protocols for data in transit, safeguarding data as it moves between client systems and storage, preventing interception or unauthorized modifications.

Third-Party Cloud Integration

- The object storage solution is designed to seamlessly integrate with major public cloud providers (e.g., AWS, Azure, GCP), enabling hybrid and multi-cloud storage environments. This allows for smooth data migration, synchronization, and management across on-premises and public cloud storage.

Load Balancing

- **Load Balancer:** Included within the solution, supporting even distribution across all nodes based on throughput requirements to ensure optimal performance and availability.

Object Uniqueness & Metadata

- **Unique Identifier:** Each stored object is assigned a unique cryptographic hash (e.g., SHA-256), along with a unique version ID, storage location, timestamp, and metadata for secure and traceable data management.

API & Protocol Support

- **Built-in S3 & REST API Compatibility:** Native support for S3, **Amazon S3** and REST API standards, eliminating the need for third-party software, hardware, or gateway for data ingestion and retrieval. This ensures seamless integration and direct interaction with storage nodes/controllers. Third party backup tools like Veeam, Comm Vault, Veritas etc. can ingest backup data to OS-100

through S3 interface.

Data Indexing & Retention

- **Indexing Support:** The solution includes native indexing capabilities, enabling efficient object search and retrieval.
- **Time-Based Retention:** Configurable retention policies at the object level, allowing administrators to manage lifecycle and retention periods based on time stamps.
- **Integration with Search Engines**
- **Search Integration:** Provides built-in integration capabilities with both internal and external search engines, facilitating advanced search and retrieval across storage environments.
- **Metadata indexing and search,** with ability to apply unlimited metadata tags/fields to the object.

Data Authenticity & Integrity

- **Hash-Based Authentication:** Utilizes robust hashing algorithms (MD5, SHA-1, SHA-2) to maintain data authenticity and integrity, ensuring secure and verifiable object storage.
- **Self-Healing and Self-monitoring Mechanism:** Automatically detects and repairs data inconsistencies to maintain data integrity locally and globally.
- **Auto-Configuration:** Dynamically configures settings across nodes to optimize performance and resilience.
- **Automated Protection & Availability Policies**
- **Customizable Availability Policies:** Ensures high availability through policies configurable for data redundancy, durability, and availability at different levels.
- **Automated Placement and Protection:** Provides automated data placement and protection strategies, which can be tailored to the organization's needs.
- OS-100 maintain transactional logs for auditing storage operations.

Replication between DC and DR Storage

- Acceleron OS-100 support replication from one S100 (DC) to another S100 (DR).

It can support unidirectional and bidirectional replication. Object storage replication is the automatic, asynchronous copying of objects (and their metadata/tags) from a source bucket to one or more destination buckets, often across different regions or accounts. As part of replication, it replicates Objects, metadata, and tags. Main use cases of replication are (a) Disaster **recovery**: Cross-region replication protects against regional outages by maintaining redundant copies in another region. (b) Compliance: to meet data retention and regulatory requirements (c) Reduced Latency: keeping data near to user reduces the data access latency (d) Data Redundancy: maintains backup copies of data in same or different regions.

Immutable and WORM

- Acceleron OS-100 support retention features like immutability and WORM. Immutable and WORM (Write once and Read Many) in object storage refer to data protection features that prevent data from being modified or deleted for a defined period.
- Once WORM enabled for specific period, data remains its pristine form until period expires. Objects are unalterable and undeletable even by administrators
- Major use cases are (a) Ransomware protection (b) Regulatory compliance like SEC Rule 17a-4(f), FINRA etc. for financial record detention

Versioning of Object

- Acceleron OS-100 support versioning of storage objects. Object versioning is a feature in object storage that keeps multiple versions of the same object in a single bucket for protecting against the accidental overwrite, deletion or corruption.
- OS100 versioning is done at bucket level and not at object level. New versions of objects are

automatically created upon new upload, deletion or overwriting of exiting object.

- OS100 maintains every object has one latest version and zero or more previous versions.
- Deleting an object adds a deletion marker and object is permanently removed only after you delete the specific version ID
- OS100 maintains versions as full objects, and it cost more storage spaces. User can retrieve any previous version directly.
- Copying the latest version to another bucket will copy only the latest version, not the previous ones.
- User cannot have versioning and bucket level retention rules simultaneously.

Data Retention

- OS100 supports retention in many ways such as snapshot retention policies, S3 life cycle rules etc.
- OS100 has built-in snapshot scheduling with retention policies that automatically prune old snapshots
- OS100 Gateway supports S3 compatible life cycle configuration for automatic object expiration and version management
- Prefix based, versioned bucket based, event based etc. life cycle management retention are supported in OS100

AD and LDAP support

- Acceleron OS-100 support LDAP and Active Directory (AD) authentication for its Object Gateway, enabling its users to authenticate to the S3 API using corporate directory credentials instead of separate Gateway credential keys.
- Only for OS100 object gateway users and not for OS100 admin or NFS clients
- OS100 supports LDAPS (LDAP over SSH) on port 636 and filter access by LDAP group membership for fine grained user authentication

Metadata Indexing and Search

- Acceleron OS100 support metadata indexing and search through elastic search integrated

to it.

- OS100 indexes object name/key, owner, size, last modified date and optionally custom metadata

IAM (Identity and Access Management)

- Acceleron OS100 supports IAM in its Object Gateway as an AWS-compatible IAM system.
- It supports accounts, users, groups, roles, and policies for multi-tenant access control
- OS100 IAM is implemented for Object Storage Gateway through S3 API.
- Multi-tenancy is achieved through tenant isolation where each account independently manages its users, groups, roles, policies and buckets.

High Availability

- Acceleron OS100 has built-in high availability features and it provide redundancy every architectural level, ensuring the cluster continues serving reads/writes even when nodes, OSDs,

network etc. fails. Acceleron OS100 works with data level redundancy techniques like data replication and/or Erasure coding (m,n).

- OS100 can achieve more than 1 node failure based on selection of erasure coding profiles (eg: (3,1), (4,2), (8,3) etc.)
- OS100 can achieve minimum high availability of 99.999% and 1 node failure with 2 disk simultaneous failure support.
- OS100 redundancy mechanism also address failure of OSD, Metadata servers, managers, network failures
- OS-100 achieves data and metadata protection using replication (default 3 copies) or erasure coding (OSD level redundancy)
- Manger level redundancy support for dashboard, monitoring level HA features
- Metadata server redundancy support for metadata and namespace level HA features
- Network redundancy and NIC bonding support for multiple public networks, cluster network failures

Specifications	
Form factor per controller	<ul style="list-style-type: none"> • 2/4U 19" rack mount • Dimensions (WxDxH) – 438 x 658 x 87 /174(mm)
Processor subsystem	<ul style="list-style-type: none"> • Intel/AMD Processors <ul style="list-style-type: none"> ○ 1/2 x AMD EPYC Processors upto 128 Core ○ 12/16/24/32 DIMM slots ○ Support DDR4/5 up to 256GB per DIMM.
Storage Type	<ul style="list-style-type: none"> • SATA 3.0 • NLSAS HDD • SAS HDD • SATA SSD • SAS SSD • NVMe SSD
Front Control	<ul style="list-style-type: none"> • Power button • System reset button
Front I/O Ports	<ul style="list-style-type: none"> • 2 x USB 3.0
Visual Indicators	<ul style="list-style-type: none"> • Power • UID • LAN activity • HDD status
Rear Panel	<ul style="list-style-type: none"> • 2 x RJ45 Network interfaces (10/100/1000 Base-T) • 1 x IPMI interface- IPMI 2.0 (10/100/1000 Base-T) • 4x25/40Gbps (Configured at purchase time) • 4x100Gbps (Configured at purchase time)
RAID Card Options	<ul style="list-style-type: none"> • 12/24Gbps with RAID 0,1,5,6,10 with 1/2/4/8GB Cache
Expansion slots	<ul style="list-style-type: none"> • 6 x FHHL PCIe 4.0 x16,
Drive bays options	<ul style="list-style-type: none"> • Configured at purchase time
Power	<ul style="list-style-type: none"> • 2 (1+1) CRPS (80+ Platinum) • 2/4U supports upto 2400W • Power Cord – C 13 , C19, 250V, 10A (India)
Cooling	<ul style="list-style-type: none"> • 80 x 25/38mm internal fans
Temperature	<ul style="list-style-type: none"> • Operating: 10°C to 35°C (50°F to 95°F) • Non-operating: -40°C to 70°C (-40°F to 158°F)
Weight	<ul style="list-style-type: none"> • 15Kg
Server Management	<ul style="list-style-type: none"> • BMC web based GUI interface to manage and monitor the sever resources.

Acceleron Labs Pvt. Ltd.
www.acceleronlabs.com

Copyright 2026

Copyright 2026 Acceleron Labs Pvt. Ltd. The information contained herein is subject to change without notice. Acceleron Labs shall not be liable for technical or editorial errors or omissions contained herein.