

Software-Defined Object Storage for Core, Edge, Hybrid Cloud Environments

Enabling Secure and Instant Access to Your Data

Key Benefits

Cost Savings

- Runs on any x86 hardware: flexibly mix and match vendors and components
- Unlimited scalability
- Streamlined administration
- Power conservation built in

Distributed Data Access

- S3/HTTP(S) access to data
- Multi-tenant capable
- Granular access controls
- Content portal for end-users

Robust Data Protection

- High availability, BC/DR
- High data durability
- Data immutability (WORM)
- Data integrity and compliance

Product Overview

DataCore Swarm provides a software-defined platform for data protection, archive, management, organization, and search at massive scale. It radically simplifies your ability to manage, store, and protect data while allowing S3/HTTP(S) access to any application, device, or end-user. Transforming your growing datasets into a flexible and immediately accessible content library, Swarm enables remote workflows and on-demand access.

With Swarm, you no longer need to migrate data into disparate storage systems for long-term preservation, delivery, and analysis. Consolidate all files, find the data you are looking for quickly, and reduce total cost of ownership.

Delivering high levels of cyber-resilience, Swarm ensures protection against security threats and compliance with business and regulatory requirements. Additionally, benefit from the ability to seamlessly shift between protection methods based on custom policies. Data governance is automated from creation to expiration—resulting in storage that adapts to your business.

Seamlessly scale from a few hundred terabytes to multiple petabytes and even exabytes with any mix of standard x86 hardware and disks. This limitless solution combines scaling of capacity and throughput with universal access and the industry's most flexible single-site and multi-site deployment options across core and edge environments.

What Users Can Expect



Automated
Data Services



Customizable
Metadata



Ransomware
Resilience



Self-Healing
Architecture



Alternative to
Tape and Cloud

Use Cases

Here are some popular use cases of Swarm catering to different industry requirements. There are numerous other applications of Swarm in organizations and service provider environments.

✓ Active Archive

Offloads data from primary NAS storage

✓ Immutable Storage for Backups

Defends against data loss and threat vectors

✓ Nearline Archive

Supports digital media workflows both in-facility and on-set (edge)

✓ Origin Storage

For OTT/VOD services and content delivery

✓ Medical Imaging Archive

Stores medical images, PACS, and VNA for healthcare sectors

✓ Archive for Digital Asset Management

Protects assets, enabling low-latency access

✓ Data Lake Storage

Handles massive workloads in research, big data, IoT, and AI/ML

✓ Multi-tenant Storage

Facilitates various cloud service offerings (e.g., SaaS)

✓ Long-term Data Preservation

Future-proofs content protection – no forklift upgrades


✓ Alternative to Public Cloud and LTO Tape

Best-suited for online, on-premises data storage

Flexible Deployment


Swarm runs on any standard x86 hardware and is highly available by design. Swarm software boots from RAM and utilizes only 5% of hard drive capacity for system data resulting in an industry-leading 95% capacity availability for your content. Scale up with disks and scale out with more nodes within a Swarm cluster, or even expand with more Swarm clusters as needed. With flexible deployment, you can add tenants and sites at any time. Swarm supports hot plug drives, adding/retiring disks/nodes and rolling upgrades of the full software stack—all with no service downtime.

Bare Metal Deployment



Boot from bare metal on any x86 server using any mix of HDDs and SSDs.

Virtual Machine Deployment



Deploy Swarm on virtual machines (VMs) powered by VMware ESXi.

Containerized Deployment



Single-server installation with all Swarm services containerized and running on Kubernetes K3s.

Packaged Hardware Appliances are also available from our partners. [Contact DataCore](#) for more information.

Swarm: Key Features

CONSUMERS			
END-USERS		APPLICATIONS & WEB SERVICES	DEVICES
ACCESS METHODS			
S3		HTTP(S)	
OPERATION & INSIGHTS	DATA SERVICES		COMMAND & CONTROL
IDENTITY & ACCESS MANAGEMENT	 WORM / IMMUTABILITY	 SYNCHRONOUS REPLICATION	WEB CONSOLE
END-USER SELF-SERVICE PORTAL	 DATA INTEGRITY SEALS	 ASYNCHRONOUS REPLICATION & DISASTER RECOVERY	REST API
AD HOC SEARCH & QUERY	 ENCRYPTION	 ERASURE CODING	AUDIT LOGS, METERING & QUOTAS
MONITORING & REPORTING	 RETENTION SCHEDULING	 SELF-HEALING	S3 OBJECT LOCK
MULTI-TENANCY	 CUSTOM METADATA	 DYNAMIC CACHING	DARKIVE™ ENERGY SAVINGS
	 UNIVERSAL NAMESPACE	 CLOUD INTEGRATION	
ANY MIX OF X86 SERVERS			
HDD		SSD	

DATA SERVICES

WORM / Immutability: Ensures data is non-erasable and non-rewritable by supporting S3 object lock and enabling legal hold. These measures protect against cyberthreats, bit rot, and other data loss risks.

Data Integrity Seals: Verifies the authenticity of stored data against tampering or corruption. This is especially helpful for auditing and compliance.

Encryption: Secures data by converting it to an unreadable format preventing unauthorized access without a key. Encryption at rest and in transit are supported.

Retention Scheduling: Manages data lifecycle with predefined time periods for retaining and deleting data based on custom policies.

Custom Metadata: Allows users to add and modify metadata (tags, descriptions, etc.) to objects, enhancing searchability, categorization, and management.

Universal Namespace: Provides a consistent view of objects across domains, buckets, and sites for easy access. Uses human-readable uniform global identifiers to retrieve objects.

Synchronous Replication: Creates local and remote real-time copies of data for immediate recovery during disruptions. This minimizes downtime and ensures data high availability.

Asynchronous Replication & Disaster Recovery: Creates redundancy across geographically dispersed locations, enabling recovery from major site outages.

Erasure Coding: Splits data into segments, encodes, and stores across different disks. This increases data availability and resilience against failures.

Self-Healing: Proactively checks for bit rot and hardware failures. Automatically rebuilds and recovers data based on data protection policies.

Dynamic Caching: Automatically persists content in RAM based on access demand (high or low), ensuring optimal infrastructure performance.

Cloud Integration: Facilitates hybrid cloud usage by copying data to and from public clouds (AWS, Azure, etc.).

NOTE: A unique feature of Swarm is that it supports both replication and erasure coding on the same node. You can set policies that shift between the two protection methods based on performance and budgetary requirements.

CONSUMERS

Entities that interface with Swarm object storage: End-users, Applications, Web Services, and Devices.

ACCESS METHODS

Protocols allowing consumers to connect with Swarm. Native access methods include S3, HTTP, and HTTPS.

Minimum Hardware Requirements

100 TB Usable Capacity with Containerized Deployment	
Processor	1 x 16-core, 2.4 GHz, Intel or AMD CPU e.g., Xeon® Silver 4314 Processor
Memory	128 GB RAM
Disk	2 x 2 TB NVMe SSDs 8 x 20 TB NL-SAS/SATA 3.5" 7200 RPM drives
Network	2 x 10 GbE
Specific requirements, such as performance, durability, and cost, influence the hardware specifications. For detailed sizing information, refer to the product documentation.	

Licensing

Simple, transparent, and flexible licensing based on usable storage capacity in TB, irrespective of the type of data stored and use case for Swarm. Opt for 1-, 3-, or 5-year terms, all including 24x7 premier support and free product updates.

OPERATION & INSIGHTS

Swarm provides a host of capabilities to simplify operations for both administrators and end-users accessing data.

- **Identity & Access Management:** Integrates with LDAP, Active Directory, and Linux PAM, as well as token-based authentication systems such as Amazon S3 API and SAML 2.0 Single Sign-On (SSO).
- **End-user Self-service Portal:** Intuitive web-based portal with distributed access for on-demand content access, sharing, streaming, video clipping, and collaboration.
- **Ad Hoc Search & Query:** Powered by Elasticsearch, Swarm enables metadata-aware search and query. Searches are dynamic, can be saved, & reused to speed up query building.
- **Monitoring & Reporting:** Monitors hardware health and cluster capacity usage, supports SNMP and Prometheus for metrics export, and integrates with Grafana for historical performance analysis.
- **Multi-tenancy:** Enables a centralized storage infrastructure to provision for distinct tenants and domains with metered quota allocation, resource management, and access control.

COMMAND & CONTROL

Swarm offers command and control through its **web console**, **REST APIs**, third-party tools, and orchestration systems. **Audit logs** provide activity tracking & compliance monitoring. **Metering** can be used to report on usage metrics for analytics and billing generation. **Quotas** are set at the object, bucket, domain, and cluster level for granular bandwidth and capacity governance. **S3 object lock** is supported to achieve data immutability and ransomware defense. **Darkive™** is a patented adaptive power conservation function that helps power down disks based on monitored periods of inactivity.

Always-on Support



Customer Excellence



10-TIME
WINNER

Customer Focused

DATAcore CSAT

99.68%

Experience round-the-clock global support from our Stevie Award-winning customer service team. Reach us via phone, email, or our online portal anytime you need assistance.

1023



Discover the Ultimate Flexibility of DataCore Software

DataCore Software delivers the industry's most flexible, intelligent, and powerful software-defined storage solutions for the core, edge, and cloud. With a comprehensive product suite, intellectual property portfolio, and unrivaled experience in storage virtualization and advanced data services, DataCore has helped over 10,000 customers worldwide modernize how they store, protect, and access data. www.datacore.com