

Swarm for HPC Environments

Hassle-Free, Multi-Tenant Object Storage Platform

BENEFITS

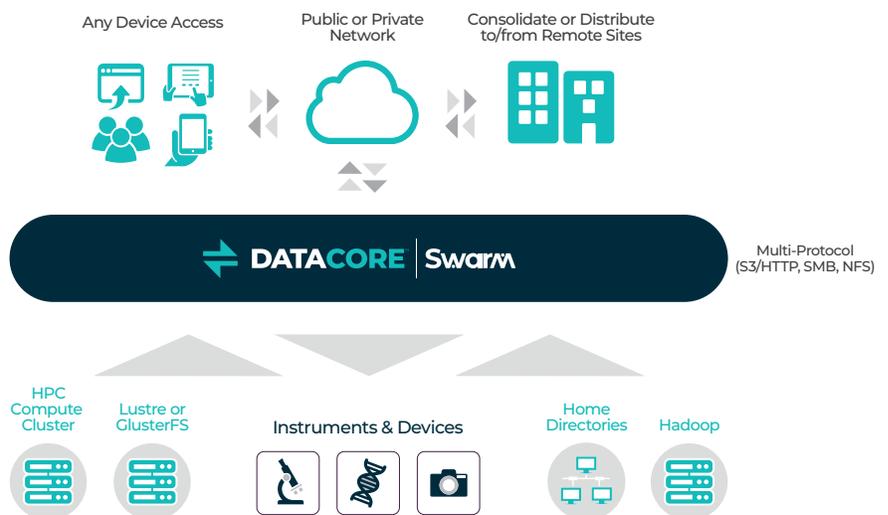
- Scale from TBs to 100s of PBs on any mix of standard hardware
- Eliminate tedious manual processes for infrastructure, data protection and tenant management
- Decrease time to discovery with integrated search and simple private/public file sharing

In high-performance computing (HPC), time to insight and discovery is determined by how quickly and easily your organization can ingest data from various sources, search for data and securely share data with authorized researchers, regardless of location. Due to performance requirements, HPC infrastructure managers have historically solved these issues using a mix of networked attached storage (NAS), distributed or parallel file systems, file transfer protocol (FTP) servers, web servers and tedious manual data authorization and access management. However, increasing capacity needs, more researchers in various locations capturing and analyzing data, along with a shift in access from traditional file protocols to RESTful interfaces are straining the status quo and hindering analysis and time to discovery.

The solution is DataCore Swarm that delivers a high-performance production infrastructure at scale simplifying content management and data access on highly resilient and cost-effective object storage. With Swarm, you no longer need to migrate data into disparate solutions for distribution, ongoing analysis and long-term preservation.

DELIVER PRIVATE OR PUBLIC STORAGE SERVICE

With Swarm you can rapidly deliver web-based storage services and provide secure access to internal or external users. Swarm snaps into existing access control systems (LDAP, AD, PAM, or token based) and comes with an integrated management portal for easy web- or API-based administration of tenants, quotas, data access and data protection policies.



RIVAL PARALLEL FILE SYSTEMS FOR READ INTENSIVE WORKFLOWS

DataCore Swarm does not require the use of any front side-caching mechanism or load balancers. Swarm's simple, flat architecture makes it low latency, self-balancing and highly symmetrical. This enables Swarm to handle many concurrent requests in parallel yielding the full throughput potential of all the drives in the system—resulting in 35 GB/s read and 12.5 GB/s write aggregate throughput via the S3 protocol on deployed HPC environments.

SIMPLE CONTENT MANAGEMENT, FILE SHARING AND SEARCH

In addition to the robust storage management user interface (UI), Swarm also has a built-in web-based UI and RESTful application program interface (API) for tenant and content management. Administrators can create buckets with capacity and bandwidth quotas and set unique data protection policies and access controls. Metadata, searches and queries are all customizable and interoperable with Elasticsearch 5 and the Elasticsearch application ecosystem. And, because the native interface to Swarm is based on HTTP, all files can be shared privately with those who have authorized access or publicly via a URL.

A BRIDGE FROM POSIX TO RESTFUL WORKFLOWS

Swarm's multi-protocol support enables interoperability with various HPC use cases making it easy to migrate data from compute clusters and network and direct attached storage (NAS and DAS) devices freeing up expensive primary storage and experimentation space. Swarm's native RESTful interface is based on HTTP 1.1 and supports the S3 protocol. The optional SwarmFS interface provides a method for S3 access to NFS data through rapid conversion of NFS data to Swarm objects delivering up to 1.6 GB/s sustained streaming in a single NFS server instance (over 3 PB per month) in HPC environments.

OPTIMIZE FILERS WHILE ELIMINATING STORAGE SILOS

In addition, DataCore offers FileFly, a Windows Server application that automates movement of cold and warm data from on-premises NAS devices and Windows filers to Swarm, Amazon S3, Google Cloud Storage or Microsoft Azure Blob Storage. FileFly enables you to consolidate data from disparate filers and tier data to the cloud for disaster recovery (DR) without disruption to end users or applications.

HPC CUSTOMERS COUNT ON DATACORE SWARM



0321



Discover the Ultimate Flexibility of DataCore Software

DataCore Software delivers the industry's most flexible, intelligent, and powerful software-defined storage solutions for block, file and object storage, helping more than 10,000 customers worldwide modernize how they store, protect, and access data. With a comprehensive product suite, intellectual property portfolio, and unrivaled experience in storage virtualization and advanced data services, DataCore is The Authority on Software-Defined Storage. www.datacore.com

GET STARTED