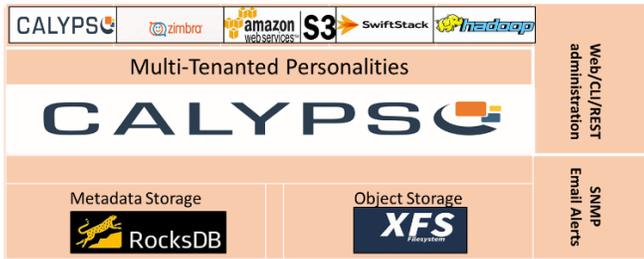


## Calypso Object Storage



### Product Essentials

#### Modular, flexible architecture

- Pluggable data storage
  - Standalone objects
  - Packed containers
- Pluggable metadata storage
  - MySQL, Couchbase, MongoDB, or Calypso
- Personalities available
  - S3, Swift, Calypso

#### Scale and performance optimized

- Billions of objects
- Hot objects caching
- Small objects
- On-line scaling nodes up or down
- Automatic Balancing

#### Reliability/Availability

- Deploy with no single points of failure
- Drive, Node, Rack, Site, Zone awareness
- Background integrity check & recovery

#### Manageability/ Security

- Applications may add own limited metadata
- Web GUI or REST APIs for management
- SNMP V2 stats and traps
- Email Alerts
- Support for ACLs, external authentication
- Configurable object expiration option
- SSL transport (Optional)

## Object Storage for Web-Scale Applications

Calypso Object Storage is a system designed to support the needs of multiple web or cloud applications running on potentially tens to hundreds of commodity storage nodes in multiple geographically distributed sites.

The Calypso developers built on their years of experience building and operating object storage systems at large multi-site service companies.

Calypso Object Storage is now available to cloud service providers and administrators of web-based applications who need a common Web-Scale object storage infrastructure that can support a multitude of different applications.

System administrators can manage and configure an easy-to-use web-based GUI or ReST-based APIs to implement object storage policies on an application by application basis.

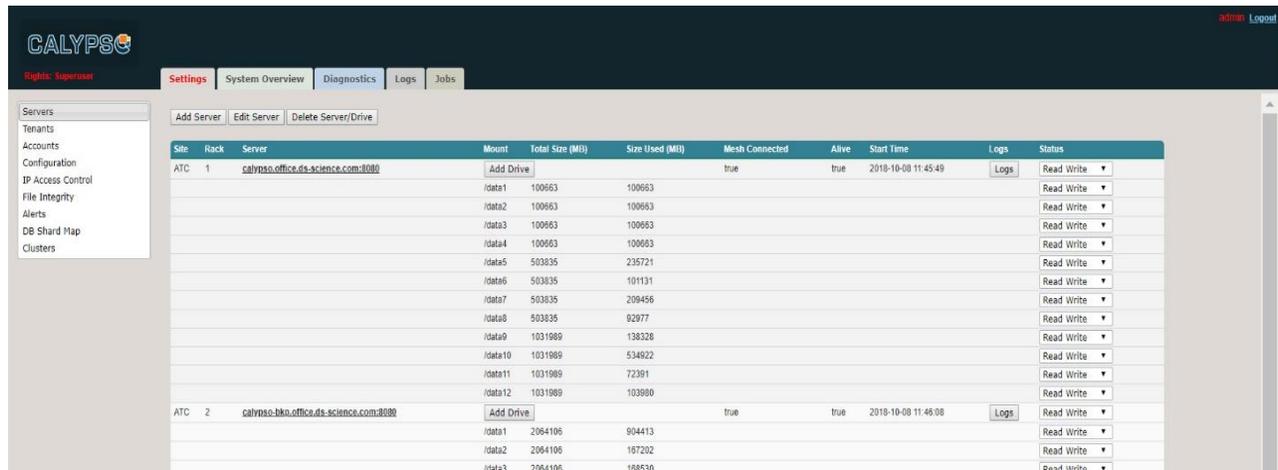
A modular architecture gives administrators flexibility to select different object protocols (Swift, S3, Calypso), eliminating the need to re-write existing applications. Administrators may also select the meta-data and data storage models that are best suited to their applications and scaling needs.

Calypso Object Storage software can be used with any storage medium including commodity host-based storage nodes.

Reliability is ensured by data replication policies that address each application's specific needs.

The system also supports policy-based data integrity checking and correction, including background, and on-access. Availability is also ensured by object copy placement across drive, server, rack and sites. The overall object storage deployment is designed to have no single points of failure.

Authentication is accomplished thru use of the built-in service or external integration with services such as Open Stack's Keystone. S3 or Swift accounts can also be created with API and secret keys. ACLs can also be implemented to control read/write access as well as access to the management console.



The screenshot shows the Calypso web interface with a navigation menu on the left and a main content area. The main content area displays a table of servers and their storage configurations. The table has columns for Site, Rack, Server, Mount, Total Size (MB), Size Used (MB), Mesh Connected, Alive, Start Time, Logs, and Status. There are two server entries, each with multiple data mounts.

Site	Rack	Server	Mount	Total Size (MB)	Size Used (MB)	Mesh Connected	Alive	Start Time	Logs	Status
ATC 1	1	calypso.office.ds-science.com:8080	Add Drive			true	true	2018-10-08 11:45:49	Logs	Read Write
			/data1	100663	100663					Read Write
			/data2	100663	100663					Read Write
			/data3	100663	100663					Read Write
			/data4	100663	100663					Read Write
			/data5	503835	235721					Read Write
			/data6	503835	101131					Read Write
			/data7	503835	209456					Read Write
			/data8	503835	92977					Read Write
			/data9	1031989	138328					Read Write
			/data10	1031989	534922					Read Write
			/data11	1031989	72391					Read Write
/data12	1031989	103980					Read Write			
ATC 2		calypso.bka.office.ds-science.com:8080	Add Drive			true	true	2018-10-08 11:45:00	Logs	Read Write
			/data1	2064106	904413					Read Write
			/data2	2064106	167202					Read Write
			/data3	2064106	168530					Read Write

## Increasing the Volume, Velocity, and Variety of Objects

Calypso Object Storage has been architected to manage and deliver the Volume, Velocity, and Variety of objects created and accessed by today's Web-scale applications.

From a **Volume** standpoint, its scale out storage nodes and modular metadata architecture is designed to support billions of objects and hundreds of petabytes. Storage nodes or drives can be added or removed on-line and the object storage system will automatically re-balance or re-copy data as needed. These processes will run in the background as to avoid causing production performance degradation.

Object access performance or **Velocity improved** over alternative approaches by setting caching policies to store frequently accessed objects or containers in multiple storage nodes. In addition, performance is improved by setting up policies to automatically place small objects in containers to avoid overhead in storing smaller objects on a file system and tracking them.

Calypso Object Storage's built-in container construct also increases the **Variety** of object

sizes; which can be stored without performance tradeoffs. Temporary objects may also have built-in expiration policies. Applications can further customize objects by adding their own meta-data and pre-fixes to the object repository database.

## Configurations

*Purchase subscription includes annual software support and maintenance contracts.*

- Software only
- Pre-configured object storage appliances (available from authorized Data Storage Science system integrators)

## Supported Storage Node Server OS

- Linux (RHEL, CentOS)
- FreeBSD 9.2+

## Data Storage Science Professional Services

- Available on a project basis.
- Services include in planning, deployment and integration of Calypso Object Storage.

## Contact Information

### Headquarters

20098 Ashbrook PI  
Suite 240  
Ashburn, VA 20147

### Sales & Support

admin@ds-science.com

### Website

www.ds-science.com