



car:ingo®

The Cloud and Object Storage Platform

CIUK 2018

Company



Pioneer in software-defined storage
Founders early tech basis of EMC Centera

Caringo surpasses
100 customers

100s of Billions of Files Stored

Recognized Leader

Unrivald Integration

Enterprise Grade

Storage Magazine
PRODUCTS OF THE YEAR 2015
FINALIST
STORAGE MAGAZINE
DUCTS OF THE YEAR 2014
FINALIST
IDC
Analyze the Future
CRN TOP 100
2014
Coolest Cloud Vendors
Gartner

VERITAS
commvault
saving forever
Telefonica
BT
netmail
ctera
DATA DEPOSIT BOX
So Simple, So Smart™

DEPARTMENT OF JUSTICE
FEDERAL BUREAU OF INVESTIGATION
Ask
St. Jude Children's Research Hospital
ALAN D. DAVY THOMAS, FOUNDER
JOHNS HOPKINS UNIVERSITY
verizon
ZirMed
Don't survive... Thrive!
TWOO

Why Object?

Scalability



Accessibility



Durability



Intelligence



Why Object in HPC?

Scale on your terms, transform archives into assets.



Internal & External
File Sharing



Built-in Tenant
Management



Rivals Parallel FS
Throughput



Scale to 100s of PBs
& Billions of Files



NFS, S3 &
SMB Support



Custom Metadata
& Search

Performance Benchmark at STFC

- Parallel file system like performance for archiving and access to warm or cold data

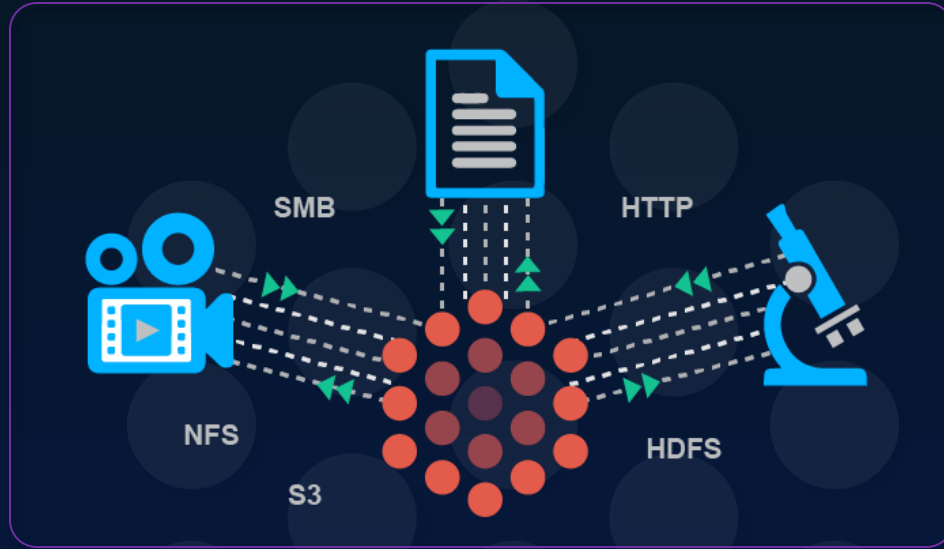
S3 Aggregate Throughput

	Performance Requirements	Caringo Results	% Above Requirement
S3 Read	21.5 GB/s	35.0 GB/s	63%
S3 Write	6.5 GB/s	12.5 GB/s	92%

NFS Single Instance Throughput

	Performance Requirements	Caringo Results	% Above Requirement
NFS Read	150 MB/s	349 MB/s	132%
NFS Write	110 MB/s	392 MB/s	256%

S3 & NFS Access



- Universal access to the same data via S3, NFS or HTTP
- Archive data using S3 APIs, read data back via NFS for simple, transparent, access from legacy tools

SwarmNFS – Centralize, Distribute, and Manage Data

Completely new approach

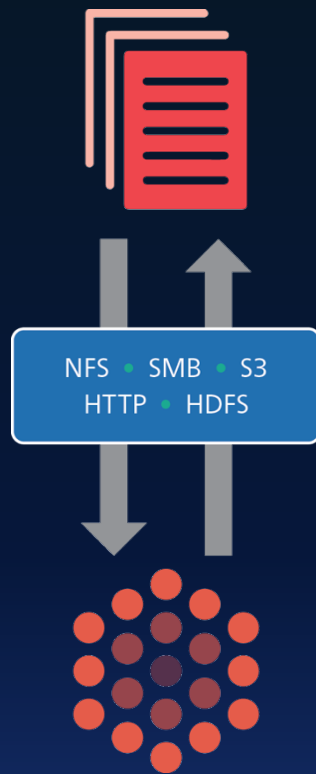
- New class of protocol conversion
 - *Not a file system gateway or connector*
- Leverages Swarm's Elasticsearch integration

Think “Object NFS”

- Every object can be written and accessed via NFS, S3 or HTTP without restriction
- Unified address space – eliminates protocol and object silos

Direct ingest and access

- Distributed mount points
- Multi-tenancy and security



Export Domains, Buckets and Collections as NFS Mounts

The screenshot displays the 'SwarmNFS Configuration' page for 'Cluster Greens_12'. It shows two NFS server groups:

- NFS SERVER GROUP 1:** Configuration URL: http://green1.tx.caringo.com:91/api/nfs/configurations/_plain/9. It contains an export named 'Marketing'.
- NFS SERVER GROUP 2:** Configuration URL: http://green1.tx.caringo.com:91/api/nfs/configurations/_plain/0. It contains two exports: 'Bear's Export' and 'Feyd-Rautha Rabban'.

Name	Scope	Domain	Bucket	Object Filter	Search Filter
Marketing	Untenanted				
Bear's Export	Tenanted	Some domain	Some buckasdasdet	immutable	Metadata
Feyd-Rautha Rabban	Untenanted			all	Metadata

The 'Add export' form includes the following fields and options:

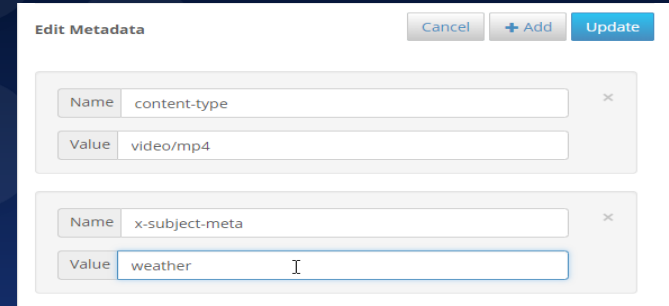
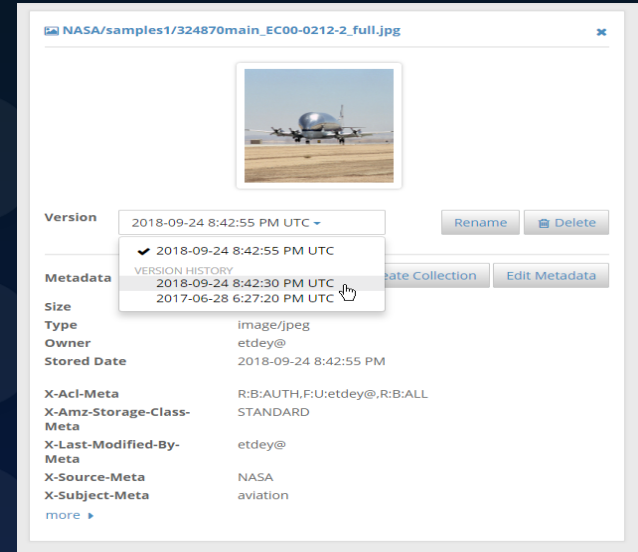
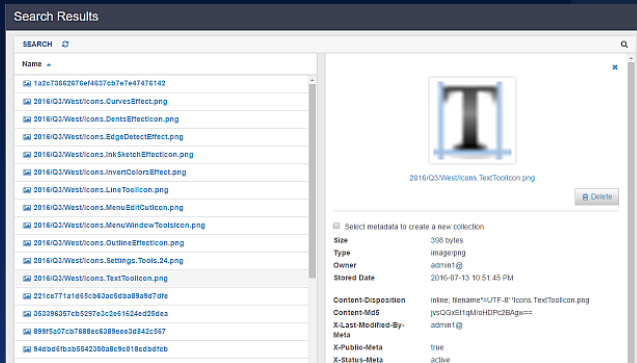
- Name:** Text input field.
- Storage IP(s) or DNS name(s):** Text input field.
- Search host(s):** Text input field.
- Scope:** Radio buttons for 'Untenanted (unnamed objects only)' and 'Tenanted (recommended)'. 'Untenanted' is selected.
- Storage interface:** Radio buttons for 'Direct to Swarm' and 'CloudScaler'. 'CloudScaler' is selected.
- FILTERS:**
 - Object type:** Dropdown menu set to 'All'.
 - Search scope:** Radio buttons for 'None', 'Metadata (new search)', and 'Collection (saved search)'. 'Metadata (new search)' is selected.
 - Q Search operator:** Dropdown menu set to 'AND'.
 - Search metadata:** A table with columns 'Name' and 'Value' for defining search criteria.
- CLOUD SECURITY:** Radio buttons for 'Access', 'Session Token', 'Single User', and 'Pass-through / None'. 'Pass-through / None' is selected.

Content Management

Provision, manage, delegate, search, organize, protect

Key Points

- Delegation and self-service with RBAC
- Collections can become NFS shares
- Data protection policy management
 - Replication/Erasure Coding
 - Versioning
- Quotas for tenant, domain, and bucket
 - Storage: logical or raw
 - Bandwidth resets at first of the month (settable)



Collections – The New Object Search Paradigm

- Catalog and organise content via standard or custom metadata strings and create “collections” based on smart views

The screenshot displays a web application interface for managing data collections. The top navigation bar shows the tenant 'abc.servic...' and domain 'data.abc.s...', with the current collection named 'Medical Data'. The user is logged in as 'admin1'.

The main interface is divided into two primary sections:

- Search Results (Left):** Shows 113 results for the 'Medical Data' collection. A search input field is present. Below, a table lists search results with columns for Name, Owner, Storage Date, Size, and Type.
- New Collection Configuration (Right):** A modal window for creating a new collection. It includes search settings and search criteria.

Search Results Table:

Name	Owner	Storage Date	Size	Type
case-0805/352238.jpg	etdey@	2016-01-19 3:54 PM	52.44 KB	image/jpeg
case-0805/352239.jpg	etdey@	2016-01-19 3:54 PM	16.51 KB	image/jpeg
case-0805/352240.jpg	etdey@	2016-01-19 3:54 PM	104.65 KB	image/jpeg
case-0805/352241.jpg	etdey@	2016-01-19 3:54 PM	101.34 KB	image/jpeg

New Collection Configuration Details:

- Search Settings:** Search Scope: Entire Domain. Buttons: Cancel, Save.
- Search Criteria:**
 - x_case_number_meta: 0805
 - x_subject_category_meta: medical
 - x_body_region_meta: neuro
- Column Headers:** Owner, Storage Date, Size, Type.


Collection Creation Options:

- Select metadata to create a new collection
- Size: 74.65 KB
- Type: image/jpeg
- Owner: etdey@
- Stored Date: 2016-01-19 3:54:15 PM
- Content-disposition: inline; filename=352243.jpg
- Content-md5: EmWXC/0h7M51zZhYfi+tdw==
- X-body-region-meta: neuro
- X-case-number-meta: 0805
- X-dc-meta-date: 2015-07-24T18:48:41
- X-dc-meta-format: image/jpeg
- X-last-modified-by-meta: etdey@
- X-source-meta: http://www.ctisus.com/teachingfiles/
- X-subject-category-meta: medical

Buttons: Create Collection, Close

Simple Internal & External Sharing

- Share objects using a global URL accessible internally or via the web (HTTP or HTTPS)
- Set object and bucket permissions using a simple access policy wizard

 <https://aoldfield.cloud.caringo.com/media/examplefile.mp4>

Access Policy [View JSON](#) [+ Add Statement](#)

Name

STATEMENT 1: CUSTOM [✎](#) [⌵](#)

PRINCIPAL

All Authenticated Users Select Users/Groups Anonymous

USERS

GROUPS

* Either a User or Group must be entered.

ACTION

Allow All Deny All Allow Select Deny Select

Allow Only Selected Allow All Except Selected

	List	Get	Put	Delete	Create	Copy	Validate	Append
Bucket <input type="text" value=""/>	<input type="checkbox"/>	<input type="checkbox"/>		<input type="checkbox"/>		<input type="checkbox"/>		
<input type="text" value="Object"/>		<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>		<input type="checkbox"/>
Object ACLs		<input type="checkbox"/>	<input type="checkbox"/>					

Multi-Tenant Management

Secure

- Connects to LDAP, Active Directory, Linux PAM
- Token-based authentication
- Coarse to fine grained user access controls

Flexible system administration

- RESTful API & via web portal
- Comprehensive naming scheme with no bucket limits
- Use your domain

Interoperable

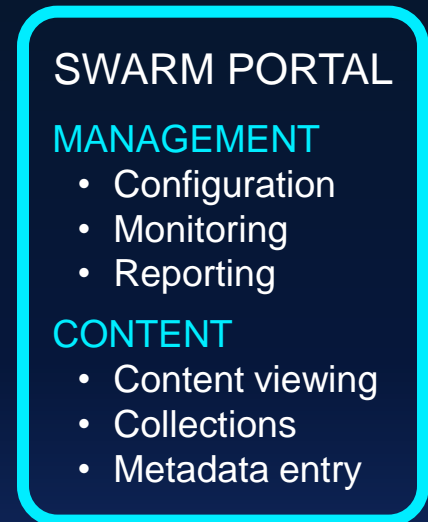
- HTTP, S3, HDFS, NFS, SMB
 - Full support for S3 versioning



Architecture Overview



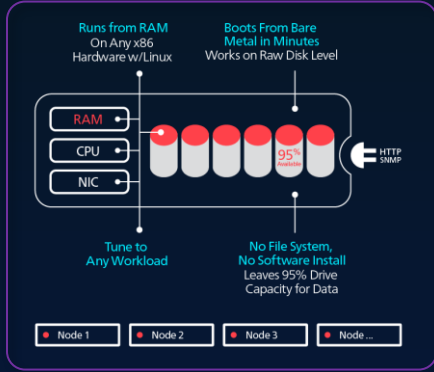
Applications | Automated Tiering from Filers



Standard x86 servers or VMs



Storage Cluster Overview




Complete Software Appliance

No file system, RAID, LUNs or DBMS required

Boots from Bare Metal;
No Additional Setup

Runs from RAM;
95% capacity for files




Flexible Deployment

Adaptive – Bio-mimicry, self-managing, market behavior

Symmetric storage nodes perform all functions

Darkive™ adaptive power conservation




Continuous Protection

EC and replicated files together (Patented)

Feeds and subclusters for replication and DR

Fast proactive recovery
AES-256 volume encryption



Instant Search and Analysis

Files contain all system and custom metadata

Unlocks data from application & location

Integrity seals,
WORM, legal hold

STFC Case Study

Geospatial data, goal to store it all

JASMIN - “super data cluster”, more bandwidth to disc than normal clusters

Benefits Achieved with Caringo


- Equivalent performance to parallel file systems for S3 throughput
- NFS and S3 access to the same objects
- Easy management of thousands of tenants and billions of objects
- Simple internal and external file access and file sharing

Kaikoura: 'Most complex quake ever studied'

By Jonathan Amos
BBC Science Correspondent

23 March 2017

[Science & Environment](#)



New wall: Whole block

The big earthquake most complex ever

Sentinel satellites to monitor every volcano

By Jonathan Amos and Rebecca Morelle
BBC science reporters

19 April 2017

[f](#) [m](#) [t](#) [e](#) [Share](#)



GETTY IMAGES

Countries that have limited resources to monitor their volcanoes will benefit most



Science & Technology
Facilities Council