Introduction to OpenStack Community Cloud Object Storage

UManitoba Spring 2026
High-Performance Computing And Cloud
Workshop

Stefano Ansaloni

University of Manitoba

May 22, 2026

University of Manitoba

Digital Research Alliance of Canada

Stefano Ansaloni

Cloud Computing Specialist at University of Manitoba
(part of the Advanced Research Computing team)

Software Developer and DevOps Specialist since 2017

Linux User/Admin since 2005

University of Manitoba

What is object storage

From Wikipedia (Object storage):

> *Object storage is a computer data storage approach that manages data as "blobs" or "objects", as opposed to other storage architectures like file systems (which manage data as a file hierarchy), and block storage (which manages data as blocks within sectors and tracks).*

University of Manitoba

Properties of object storage

- Retention of massive amounts of unstructured data
- Data is written once and read once (or many) times
- Each object is typically associated with a variable amount of metadata, and a globally unique identifier
- Not intended for transactional data
- No locking and sharing mechanisms needed to maintain a single, accurately updated version of a file

Wikipedia [Object storage]

University of Manitoba

OpenStack object storage

OpenStack supports object storage through its component *Swift*.

Swift is responsible for ensuring data replication and integrity across the OpenStack infrastructure.

Swift also implements a subset of the Amazon S3 operations, making it compatible with most of the S3 clients/libraries.

OpenStack [Introduction to Object Storage]
OpenStack [S3/Swift REST API Comparison Matrix]

University of Manitoba

Arbutus object storage

All Arbutus projects are allocated a default 1TB of object storage space (up to 10TB for RAS).

More than 10TB must be requested and allocated under the annual RAC.

Management of a project's object storage containers/buckets is self-service (see cloud storage options).

DRAC Wiki [Arbutus object storage]

University of Manitoba

Accessing Arbutus object storage

Arbutus offers access via two different protocols: Swift or Amazon S3.

There are several ways to access data containers/buckets:

- using an S3-compatible client (e.g. *s3cmd*)
- using Globus
- If the policy is set to "public", using a browser via an HTTPS endpoint:
  `https://object-arbutus.alliancecan.ca/PROJECT_ID:DATA_CONTAINER/FILENAME`
  (legacy: `https://object-arbutus.cloud.computecanada.ca/DATA_CONTAINER/FILENAME`)

University of Manitoba

DRAC Wiki [Arbutus object storage]

Creating a container/bucket on Arbutus

To create a new container/bucket, use the left menu to select
Project » Object Store » Containers » +Container button

A form will display, asking to define the container properties.

In the *Container Name* field you must choose a name for the container.

Using the *Container Access* options, it is possible to make it publicly available.

University of Manitoba

Configuring s3cmd for Arbutus object storage

1. Install the OpenStack command line client
2. Install *s3cmd*
3. Generate the access and secret keys using the OpenStack command line client with "`openstack ec2 credentials create`"
4. Configure *s3cmd* with "`s3cmd --configure`":
   - `Access Key: 32_DIGIT_ACCESS_KEY`
   - `Secret Key: 32_DIGIT_SECRET_KEY`
   - `Default Region []: US`
   - `S3 Endpoint []: object-arbutus.alliancecan.ca`
     (legacy: `object-arbutus.cloud.computecanada.ca`)
   - `DNS-style bucket+hostname:port template for accessing a bucket []:`
     `object-arbutus.alliancecan.ca`
     (legacy: `object-arbutus.cloud.computecanada.ca`)

University of Manitoba

DRAC Wiki [Accessing object storage with s3cmd]

S3cmd basic commands

Create container/bucket	`s3cmd mb s3://BUCKET_NAME/`
Show container/bucket status	`s3cmd info s3://BUCKET_NAME/`
Upload files	`s3cmd put --guess-mime-type FILE_NAME.dat \` `s3://BUCKET_NAME/FILE_NAME.dat`
Delete files	`s3cmd rm s3://BUCKET_NAME/FILE_NAME.dat`
Set ACLs	`s3cmd setacl --acl-public -r s3://BUCKET_NAME/` `s3cmd setacl --acl-private -r s3://BUCKET_NAME/`

University of Manitoba

Demo

Serving static web content through
Arbutus object storage

University of Manitoba

Serving static web content through Arbutus object storage

Specifications

Cloud system $\Rightarrow$ Arbutus cloud

Content $\Rightarrow$ Potree – https://github.com/potree/potree
(WebGL based point cloud renderer for large point clouds)

University of Manitoba

Final thoughts

The information presented (along with the demonstration) are meant to be as generic as possible, in order to provide a basic knowledge of cloud systems and OpenStack functioning.

Useful links:

- DRAC cloud wiki page
- DRAC cloud storage options wiki page
- DRAC object storage wiki page
- DRAC object storage clients wiki page

Remember to **read and understand** your responsibilities.

University of Manitoba

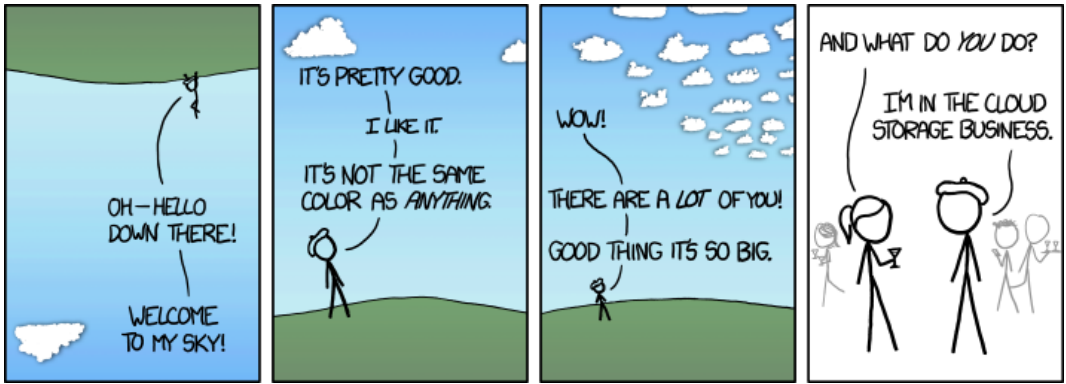

XKCD [1117]

University of Manitoba

Questions?

University of Manitoba

Thank you

University of Manitoba