




Science and
Technology
Facilities Council



iris

IRIS DynaFed: IAM-Integrated Echo Storage

IAM Users Workshop
January 28th 2021
Sam Glendenning



Today's checklist



The Echo Storage service



DynaFed



IRIS DynaFed – a relaunch!



Demo



Features & Use Cases



Echo S3 Data Storage

- Maintained by Data Services team
- Online, scalable, fast object storage service accessible from anywhere on the internet
- Echo is STFC's instantiation of this and it provides Amazon's S3 API
- High data throughput and capacity
- Perfect for an age of increasing data size and demand

- User accounts are created which can then create object storage spaces called **buckets**
- Account assigned an access and secret key, used to interact with buckets and their contents



Object Storage

```
GET /object/Kbg1Bn7qepo  
PUT /object/Kbg1Bn7qepo
```

Object Storage

Echo S3 Data Storage

Great because...

- Object store scales more regardless of file size (performance is generally constant)
- Handy for storing large amounts of data long-term
 - Logs
 - Scientific Datasets
 - Experiment Results
 - Archives
 - etc.
- Can share files with presigned URLs (time-limited)
- Meets a storage need for lots of communities

Limited by...

- Sharing of data restricted to presigned URLs or sharing keys
- Keys can be misplaced, accidentally uploaded with source code, etc.
- Buckets are designed to be private and setting up multi-user access can be difficult
- Typical interaction through software, scripts or the command line











```
https://s3.echo.stfc.ac.uk/group/bucket/path/to  
/file?AWSAccessKeyId=92E52RHJJMOBADPBUH5  
H&Signature=j987ZfNuDcVPVbc9kEvE72qapNr%3  
D&Expires=1606488466
```

- URL of object
- Access key for this bucket
- The unique random signature of verification
- Time of expiry (in 2 hours)

Example pre-signed URL

DynaFed

- Federated storage service that can act as a gateway to Echo S3 buckets
- Provided through a browser interface
- Simulates a file system look and feel for familiarity and ease of navigation
- Manages bucket keys and provides a layer of user authentication:
 - DynaFed stores the keys to a bucket
 - You prove who you are to DynaFed **and** that you are allowed to access this bucket
 - DynaFed provides access to the bucket

Mode	Links	UID	GID	Size	Modified	Name
drwxrwxrwx	0	0	0	0	Thu, 01 Jan 1970 00:00:00 GMT	 atlas
drwxrwxrwx	0	0	0	0	Thu, 01 Jan 1970 00:00:00 GMT	 dmc
drwxrwxrwx	0	0	0	0	Thu, 01 Jan 1970 00:00:00 GMT	 dteam
drwxrwxrwx	0	0	0	0	Thu, 01 Jan 1970 00:00:00 GMT	 dune
drwxrwxrwx	0	0	0	0	Thu, 01 Jan 1970 00:00:00 GMT	 enmr
drwxrwxrwx	0	0	0	0	Thu, 01 Jan 1970 00:00:00 GMT	 lhcb
drwxrwxrwx	0	0	0	0	Thu, 01 Jan 1970 00:00:00 GMT	 ligo
drwxrwxrwx	0	0	0	0	Thu, 01 Jan 1970 00:00:00 GMT	 prominence
drwxrwxrwx	0	0	0	0	Thu, 01 Jan 1970 00:00:00 GMT	 rucioecdf
drwxrwxrwx	0	0	0	0	Thu, 01 Jan 1970 00:00:00 GMT	 ska

DynaFed

Right now...

- Authorises users based on an x.509 certificate
- Requires manual editing of authorisation files (adding and removing certificates)
- Not widely used, x.509 certificates can be a hassle
- Quite limited web interface

With new release...

- Opening up the platform for IRIS communities
- Expanding on its use cases
- More dynamic and easier to use
- Intuitive and more expansive user authentication

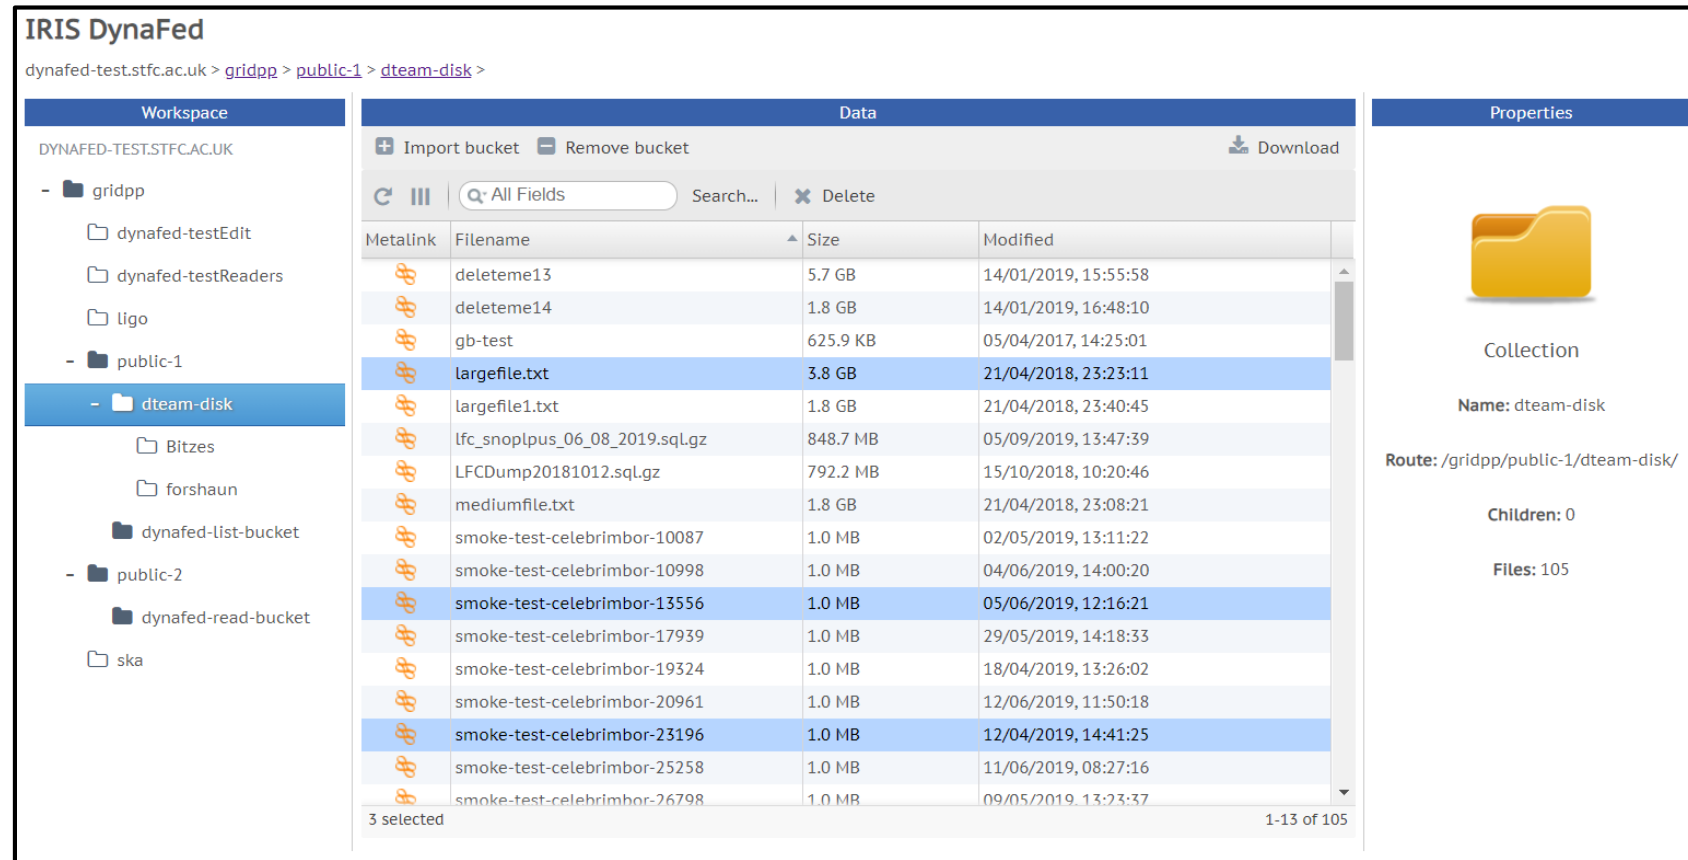


IRIS DynaFed – a relaunch!

- Integrated with the IRIS IAM authentication system
- Uses OpenIDC to grab info about the user to authenticate them
- Uses IAM groups to grant the user access to buckets

A new system:

If the user is a member of Group A on IAM, grant them access to Group A's buckets on DynaFed



The screenshot displays the IRIS DynaFed web interface. The breadcrumb path is `dynafed-test.stfc.ac.uk > gridpp > public-1 > dteam-disk >`. The interface is divided into three main sections: Workspace, Data, and Properties.

Workspace: Shows a tree view of the file system. The current path is `public-1 > dteam-disk`. Other visible folders include `Bitzes`, `forshaun`, `dynafed-list-bucket`, `public-2`, `dynafed-read-bucket`, and `ska`.

Data: A table listing files and folders. The table has columns for `Metalink`, `Filename`, `Size`, and `Modified`. The files listed are:

Metalink	Filename	Size	Modified
	deleteme13	5.7 GB	14/01/2019, 15:55:58
	deleteme14	1.8 GB	14/01/2019, 16:48:10
	gb-test	625.9 KB	05/04/2017, 14:25:01
	largefile.txt	3.8 GB	21/04/2018, 23:23:11
	largefile1.txt	1.8 GB	21/04/2018, 23:40:45
	lfc_snoplus_06_08_2019.sql.gz	848.7 MB	05/09/2019, 13:47:39
	LFCDump20181012.sql.gz	792.2 MB	15/10/2018, 10:20:46
	mediumfile.txt	1.8 GB	21/04/2018, 23:08:21
	smoke-test-celebrimbor-10087	1.0 MB	02/05/2019, 13:11:22
	smoke-test-celebrimbor-10998	1.0 MB	04/06/2019, 14:00:20
	smoke-test-celebrimbor-13556	1.0 MB	05/06/2019, 12:16:21
	smoke-test-celebrimbor-17939	1.0 MB	29/05/2019, 14:18:33
	smoke-test-celebrimbor-19324	1.0 MB	18/04/2019, 13:26:02
	smoke-test-celebrimbor-20961	1.0 MB	12/06/2019, 11:50:18
	smoke-test-celebrimbor-23196	1.0 MB	12/04/2019, 14:41:25
	smoke-test-celebrimbor-25258	1.0 MB	11/06/2019, 08:27:16
	smoke-test-celebrimbor-26798	1.0 MB	09/05/2019, 13:23:37

At the bottom of the table, it indicates `3 selected` and `1-13 of 105` items.

Properties: Shows a folder icon and the following details:

- Collection
- Name: dteam-disk
- Route: /gridpp/public-1/dteam-disk/
- Children: 0
- Files: 105

IAM and DynaFed in tandem

IAM:

- Set up groups for your team/purpose
- Join groups by sending join requests
- Accept or deny requests to join your own groups

DynaFed:

- Import and remove Echo buckets and assign them to your groups
- Browse other buckets imported by your team mates

Sam Glendenning



Sam Glendenning

dgz17987

b09fc4dc4-db16-493a-9bb8-9733f623f692

Email sam.glendenning@stfc.ac.uk

Status ✓ Active

Groups

dynafed/testEdit

dynafed/testReaders

Group requests

No request found

➔ Join a group

IAM Dashboard

IRIS DynaFed

dynafed-test.stfc.ac.uk > [gridpp](#) > [dynafed-testReaders](#) > [dteam-disk](#) >

Workspace

DYNAFED-TEST.STFC.AC.UK

- gridpp

- dynafed-testEdit

dteam-disk

ligo-frames

samglendenning-testbuc

- dynafed-testReaders

- dteam-disk

Bitzes

forshaun

dynafed-read-bucket

+ Import bucket - Re

🔄 ||| 🔍 All Fields

Metalink	Filename
🔗	deleteme13
🔗	deleteme14
🔗	gb-test
🔗	largefile.txt
🔗	largefile1.txt
🔗	lfc_snoplpus_06_
🔗	LFCDump201810
🔗	mediumfile.txt
🔗	smoke-test-celet
🔗	smoke-test-celet

DynaFed homepage

Demo

Feel free to try for yourself. You can find it at:

<https://dynafed-test.stfc.ac.uk>

Apply for an IAM account at:

<https://iris-iam.stfc.ac.uk/login>

IAM help at:

<https://iris-iam.stfc.ac.uk/help/>

The screenshot shows the IRIS DynaFed interface. The top navigation bar includes the workspace name 'DYNAFED-TEST.STFC.AC.UK' and the current path 'gridpp'. Below this is a table with columns for 'Met...', 'Filename', 'Size', and 'Modified'. To the right of the table is a 'Properties' panel for the 'gridpp' collection, showing details such as 'Name: gridpp', 'Route: /gridpp/', 'Children: 6', and 'Files: 0'. The interface also features buttons for 'Import bucket', 'Remove bucket', and 'Download'.

The screenshot shows the IRIS IAM user profile page for Sam Glendenning. The page is divided into several sections: a header with the user's name and email, a 'Groups' section listing 'dynafed/testEdit' and 'dynafed/testReaders', a 'Group requests' section with a 'Join a group' button, and a 'Linked accounts' section with a 'Link external account' button. The user's status is 'Active', and the page includes an 'Edit Details' button at the bottom.

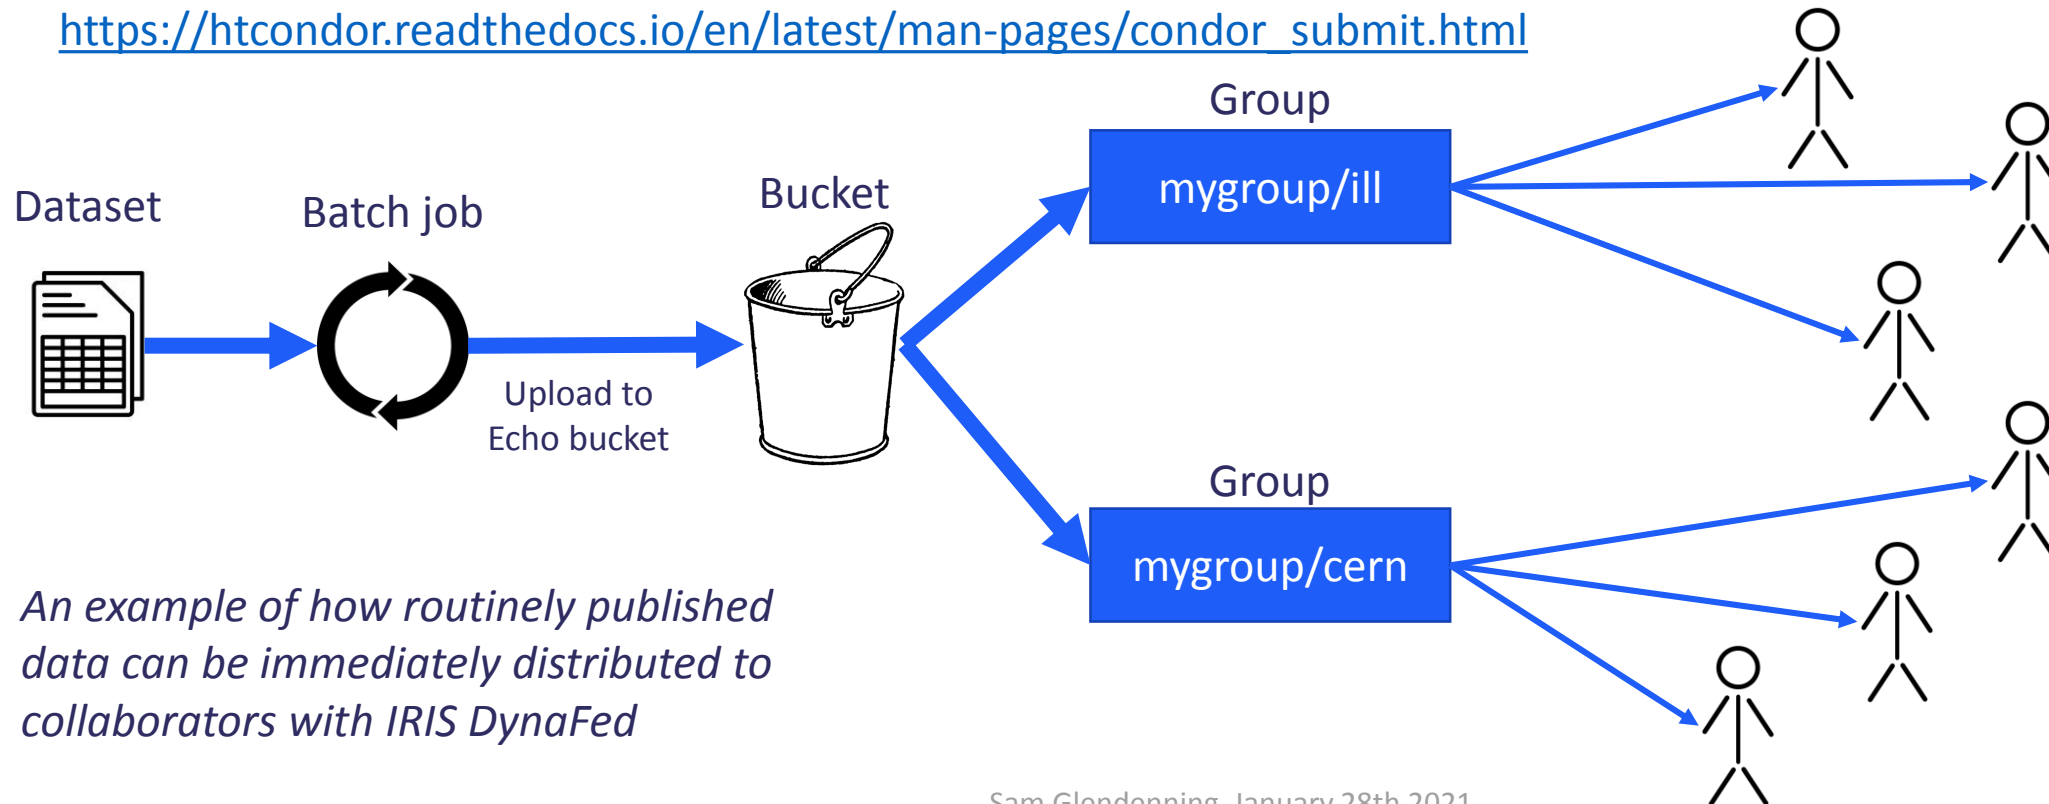
Features

- Straightforward sharing of files with other users, both time-limited and not
- General purpose users don't require knowledge of bucket keys (unless importing/removing buckets)
- Bulk downloading of files
- No software required
- Simple, intuitive web interface
- Access to buckets without needing to be on a VPN, internal network, etc.





Uploading to Echo buckets

- **Python boto3 library** - Great for use in custom scripts and jobs that can be run quickly, either manually or automatically, to upload files <https://boto3.amazonaws.com/v1/documentation/api/latest/guide/quickstart.html>
- **s3cmd** - Useful for getting started and configuring buckets, offering a wide array of functions for uploading and downloading data <https://s3tools.org/s3cmd-howto>
- **rclone** – Great for syncing and managing files on cloud storage, including S3. Allows for multi-threaded jobs <https://rclone.org/>
- **HTCondor** - Can presign URLs and perform time-limited uploads without exposing access credentials https://htcondor.readthedocs.io/en/latest/man-pages/condor_submit.html



An example of how routinely published data can be immediately distributed to collaborators with IRIS DynaFed



Use Cases

- Collaboration between users working in the same team or for a similar goal
- Multi-tiered access to various buckets (one group has access to some buckets, another has access to others)
- Regular publishing of data to collaborators at various research institutes
- Sharing relevant files with users who don't have an IRIS IAM account without revealing any sensitive information or having to use the IRIS DynaFed platform

Try now at:

<https://dynafed-test.stfc.ac.uk>



Science and
Technology
Facilities Council

Going live soon at:

<https://dynafed.stfc.ac.uk>



Questions? Discussions?

sam.glendenning@stfc.ac.uk



Science and
Technology
Facilities Council

Thank you



Science and Technology Facilities Council



@STFC_matters



Science and Technology Facilities Council