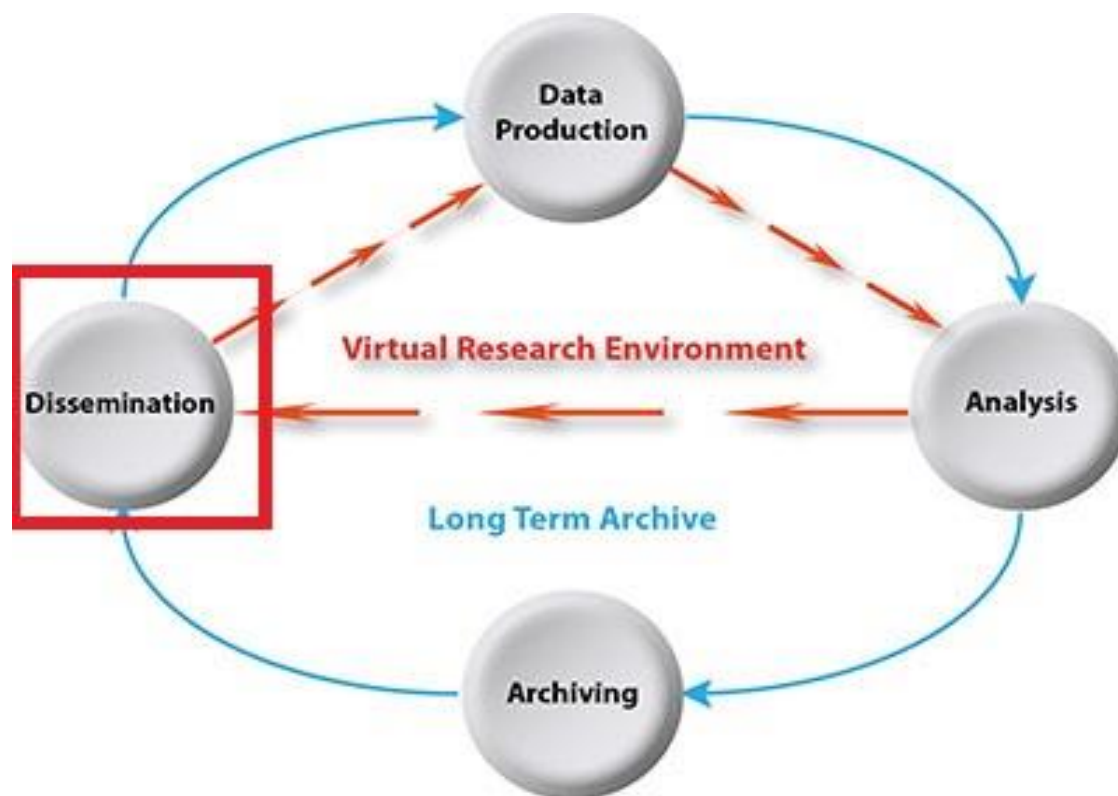


Cloud-Storage @ DKRZ

An introduction to Openstack Swift and it's API

Christopher Bartz
Deutsches Klimarechenzentrum (DKRZ)

Workflow



Goal

- Access data from different devices, e.g. notebook, smartphone and supercomputer
- Share data with colleagues with or without account at DKRZ
- Replicate data between different sites automatically

Openstack Swift - Characteristics

- Geo-distributed objectstore
- Horizontal scalable
- Software-defined storage
- Eventual consistency
- RESTful API



Openstack Swift – Pros & Cons

- + World-wide accessible
- + Easy sharing with externals
- + Simple RESTful API via HTTP
- + Site-replication
- + Filesystem independence

- No filesystem features (e.g. no renaming)
- Only atomic object manipulation

Openstack Swift - Features

- Temporary URLs
- Bulk operations
- Object expiration
- ACLs
- Replication (Container/Cluster)
- Versioning
- Metadata
- Static/Dynamic large objects
- Static Web

Concepts

- Accounts: Top-Level of hierarchy
- Containers: Namespace for objects
- Objects: Store data & metadata

Interfaces





- RESTful
- Web
- Command-line
- Python

Swiftbrowser (developed @ DKRZ)

Swiftbrowser

k20200:k202103 ▾

Containers

Name	Objects	Size	
 dkrz_scratch	2	1.6 KB	
 tempcontainer	3527	16.3 GB	

16.3 GB of 143.0 GB used

Python

```
from swiftclient.client import Connection

c = Connection(...) # Credentials

c.put_container("my_container")

metadata, objects = c.get_container("my_container")

f = open("file")
c.put_object("my_container", "my_object", f)

obj = c.get_object("my_container", "my_object")
```

Command-line

```
$ swift post my_container
```

```
$ swift list my_container
```

```
$ swift upload my_container file
```

```
$ swift download my_container file
```

RESTful

```
curl -XPUT -H 'X-Auth-Token: token'
```

```
https://swift.dkrz.de/v1/dkrz\_account/my\_container
```

```
curl -XGET -H 'X-Auth-Token: token'
```

```
https://swift.dkrz.de/v1/dkrz\_account/my\_container
```

```
curl -XPUT --data-binary @file -H 'X-Auth-Token: token'
```

```
https://swift.dkrz.de/v1/dkrz\_account/my\_container/my\_object
```

```
curl -XGET -H 'X-Auth-Token: token'
```

```
https://swift.dkrz.de/v1/dkrz\_account/my\_container/my\_object
```

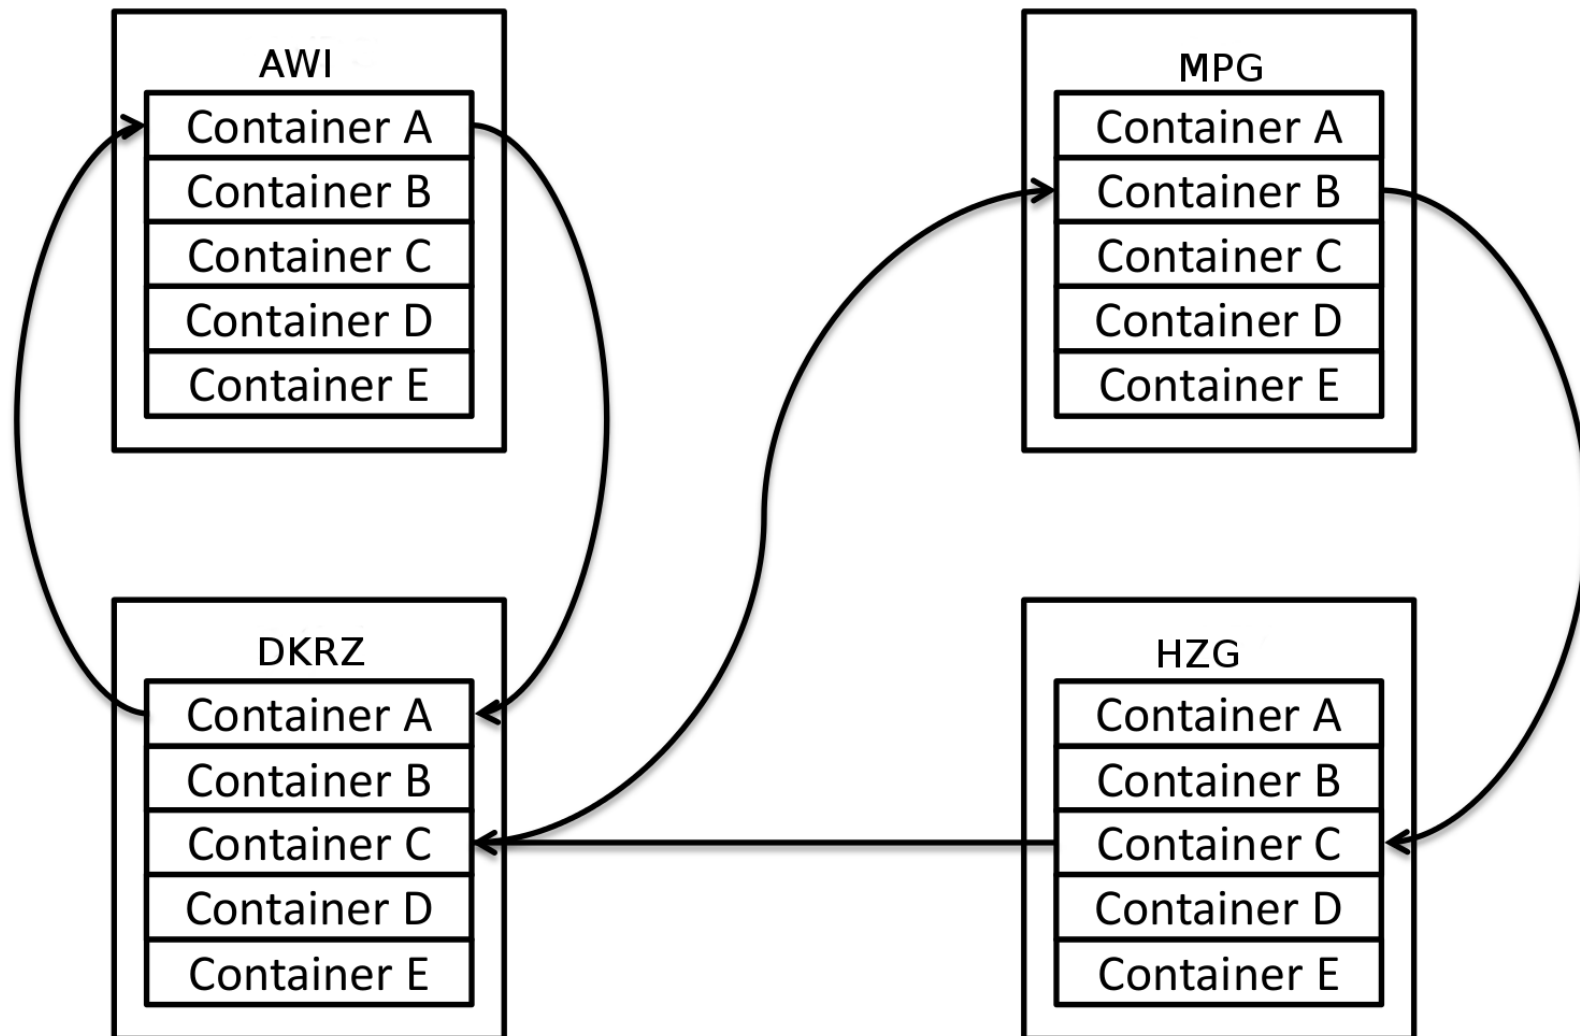
Use case - Possible workflow

- Generate data.
- Do postprocessing.
- Copy data with python client into the cloud.
- Send links to external people or retrieve data at home via swiftbrowser.

Use case – Build an application

- Swift as backend-store for static data (backup, media,....)
- World-wide accessible
- Filesystem-independency
- Clear and easy API

Use case - Container Synchronization



Swift @ DKRZ

- **Capacity:** 300 TB brutto
- No encryption!
- No backup!
- **Accounts:**
 - Private space (5 GB).
 - Each HLRE project (0 to 20 TB).
- **Scratch:** 5 TB each account.

Useful links

- Cloud @ DKRZ: <https://www.dkrz.de/Nutzerportal-en/doku/cloud>
- Building Applications with Swift: <https://youtu.be/rAGjrquuoC8>
- Python Swiftclient: <http://docs.openstack.org/developer/python-swiftclient/>
- Swift in general: <http://docs.openstack.org/developer/swift/>
- RESTful API: <http://developer.openstack.org/api-ref-objectstorage-v1.html>