



## H3: An embedded object store

Antony Chazapis

chazapis\_at\_ics.forth.gr

FOSSCOMM 2020, 20-21/11/2020

# Object stores

Most popular type of **storage-as-a-service**

**Buckets** that hold **objects** with no hierarchy

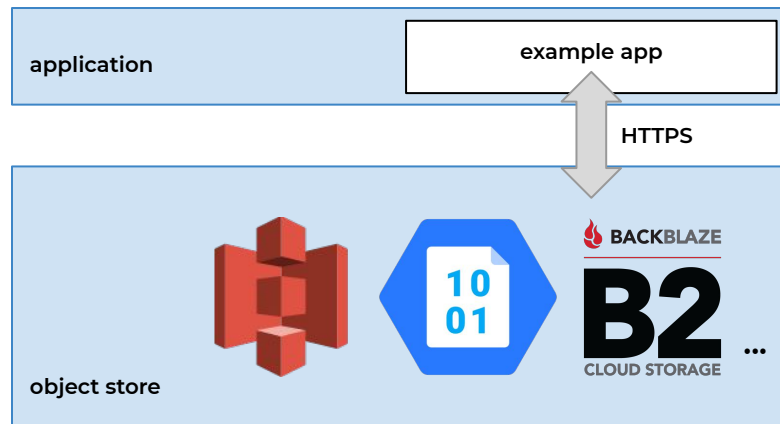
Pros and cons

- + No management
- + Unlimited capacity
- + Wide range of offerings
- Cost
- Performance

Local solutions (prototyping, testing, etc.)

MinIO

s3proxy



# H3 (h3lib, h3cli, h3fuse)

H3 is an **embedded object store (library)**

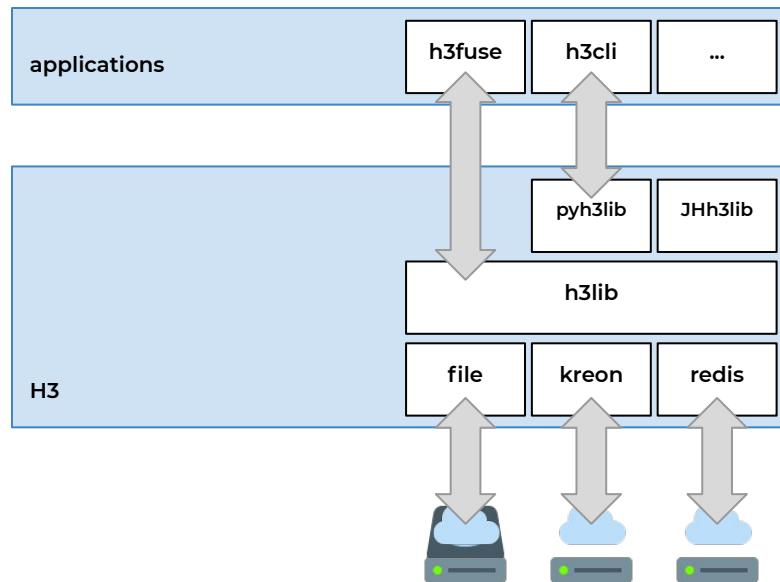
API calls for objects **converted** to KV operations implemented by plug-ins

**Plug-ins** for filesystem, Kreon, Redis, RocksDB

**Clean API** with Python and Java wrappers

**CLI** for management

**FUSE-based filesystem** for file semantics



# Usage

Simple API (bucket, object, multipart)

Python, Java, C available

Python

```
from pyh3lib import H3

h3 = H3('redis://127.0.0.1:6379')
h3.create_bucket('mybucket')
h3.create_object('mybucket', 'a.txt', data)

h3.list_objects('mybucket') # Returns ['a.txt']
```

Bash

```
# h3cli --storage "redis://127.0.0.1:6379" mb h3://mybucket
# h3cli --storage "redis://127.0.0.1:6379" cp a.txt
h3://mybucket/a.txt
# h3cli --storage "redis://127.0.0.1:6379" ls h3://mybucket
a.txt
#
```

# Design

Runs as part of the application (no HTTP)

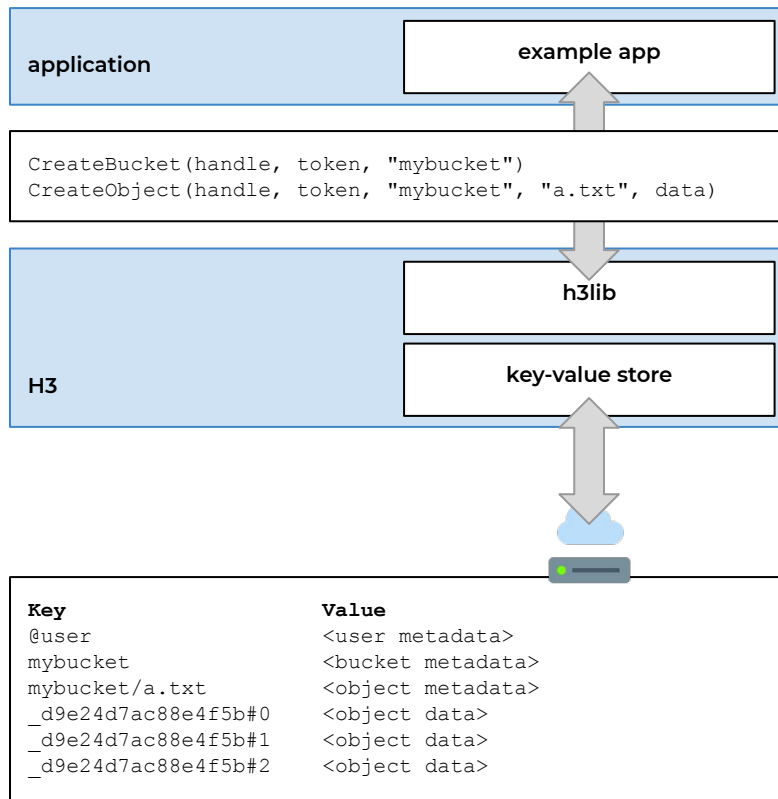
Data and metadata stored as key-values

Extensible backend API

Selection of storage methods

Simpler, more lightweight calls

May exploit high-speed networks



# Design

Runs as part of the application (no HTTP)

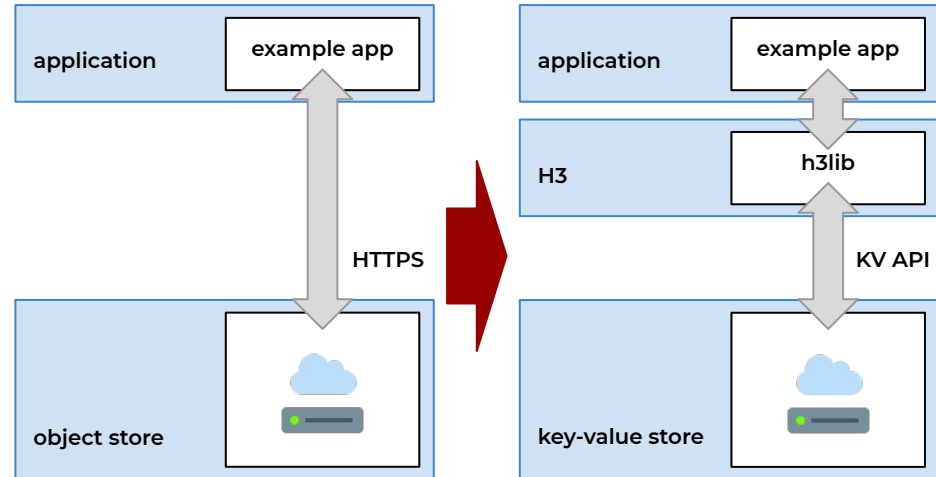
Data and metadata stored as key-values

Extensible backend API

- Selection of storage methods

- Simpler, more lightweight calls

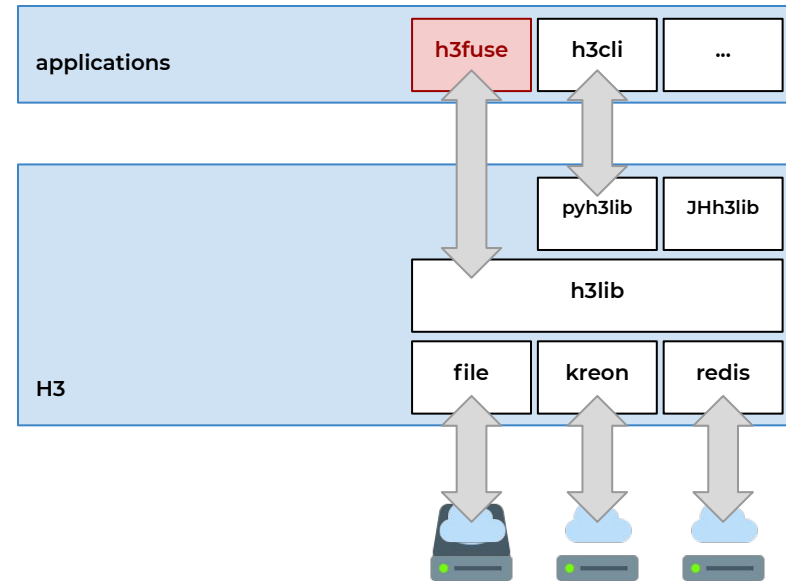
- May exploit high-speed networks



# h3fuse

Files for compatibility

Kubernetes integration




# csi-h3

Implementation of a **Container Storage Interface (CSI)** plugin

Easy provisioning of storage to Kubernetes containers via **h3fuse**

**Single** controller and nodeplugin

Based on available **H3** container





# csi-h3 flow

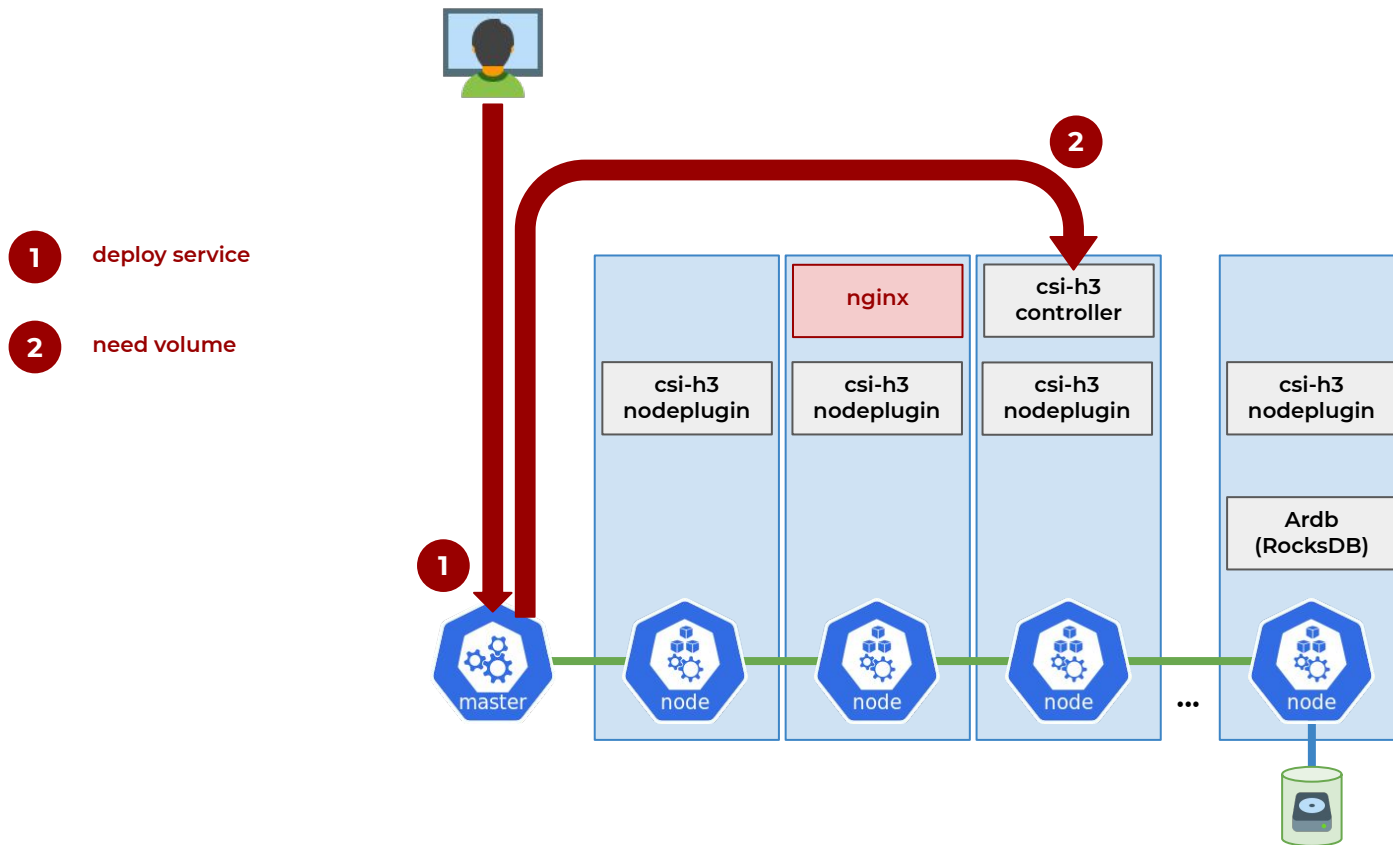
1 deploy service




```

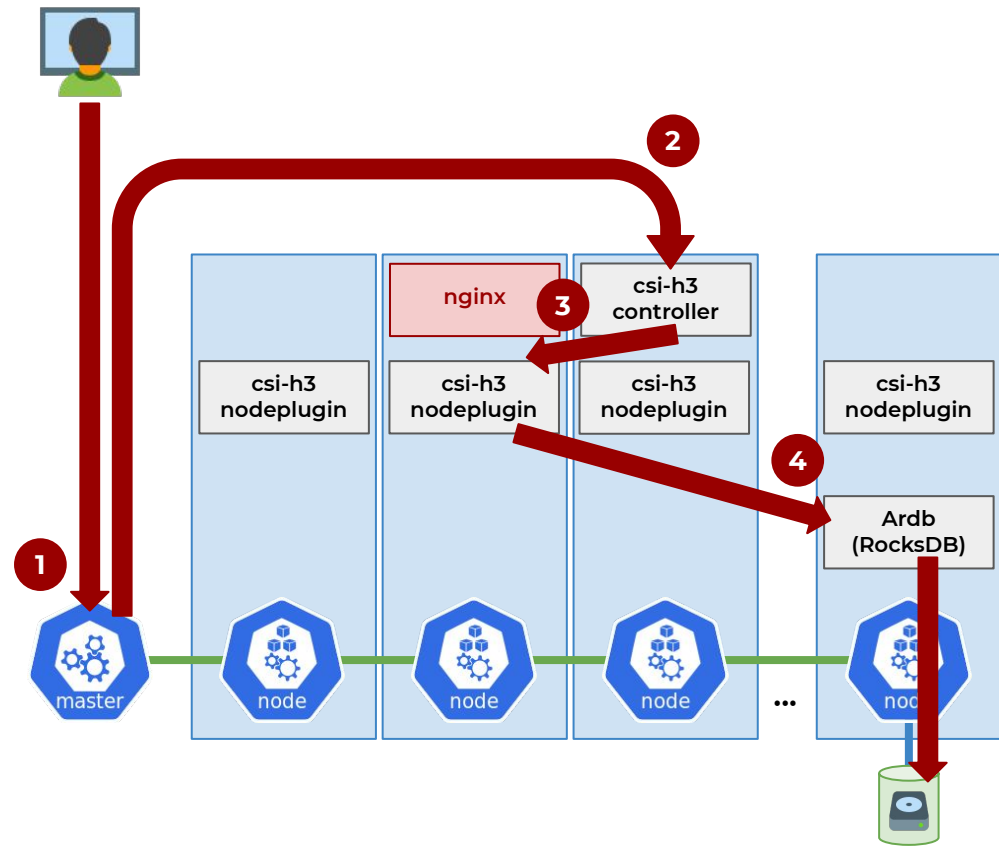
apiVersion: v1
kind: PersistentVolume
metadata:
  name: h3-example
spec:
  accessModes:
  - ReadWriteMany
  capacity:
    storage: 10Gi
  storageClassName: h3
  csi:
    driver: csi-h3
    volumeHandle: h3-example
    volumeAttributes:
      storageUri: "redis://redis.default.svc:6379"
      bucket: "mybucket"
---
apiVersion: v1
kind: PersistentVolumeClaim
...
---
apiVersion: v1
kind: Pod
...
spec:
  ...
  volumes:
  - name: h3-example
    persistentVolumeClaim:
      claimName: h3-example

```

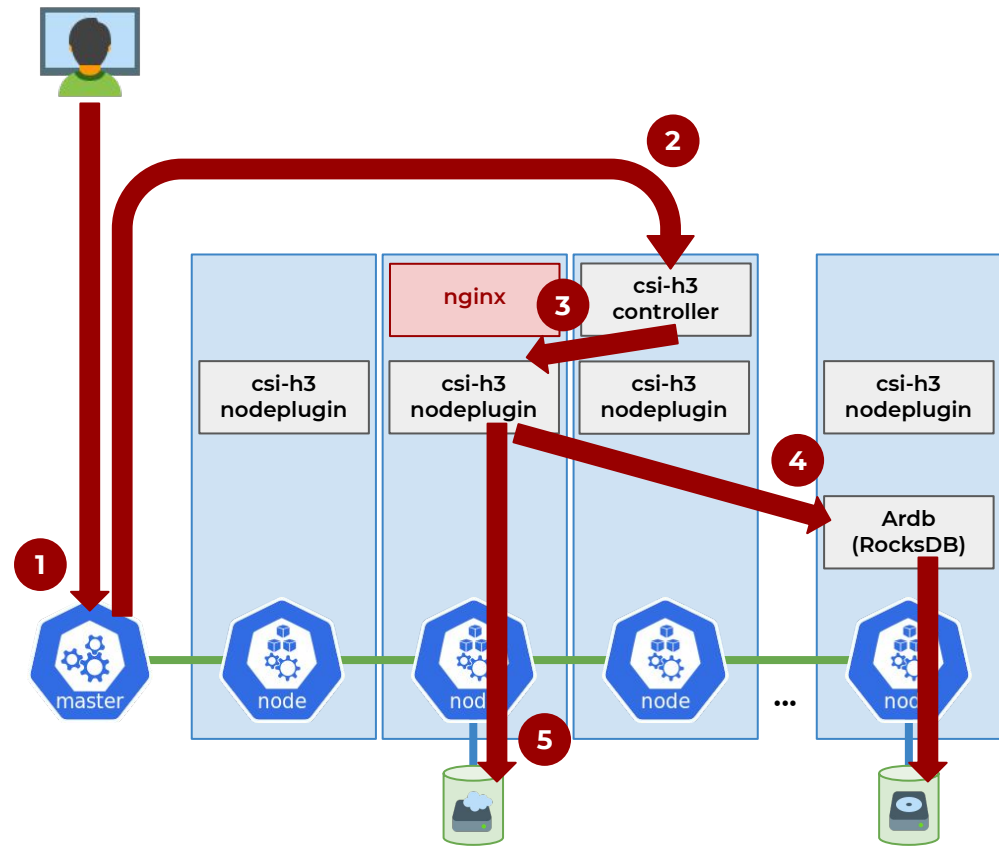




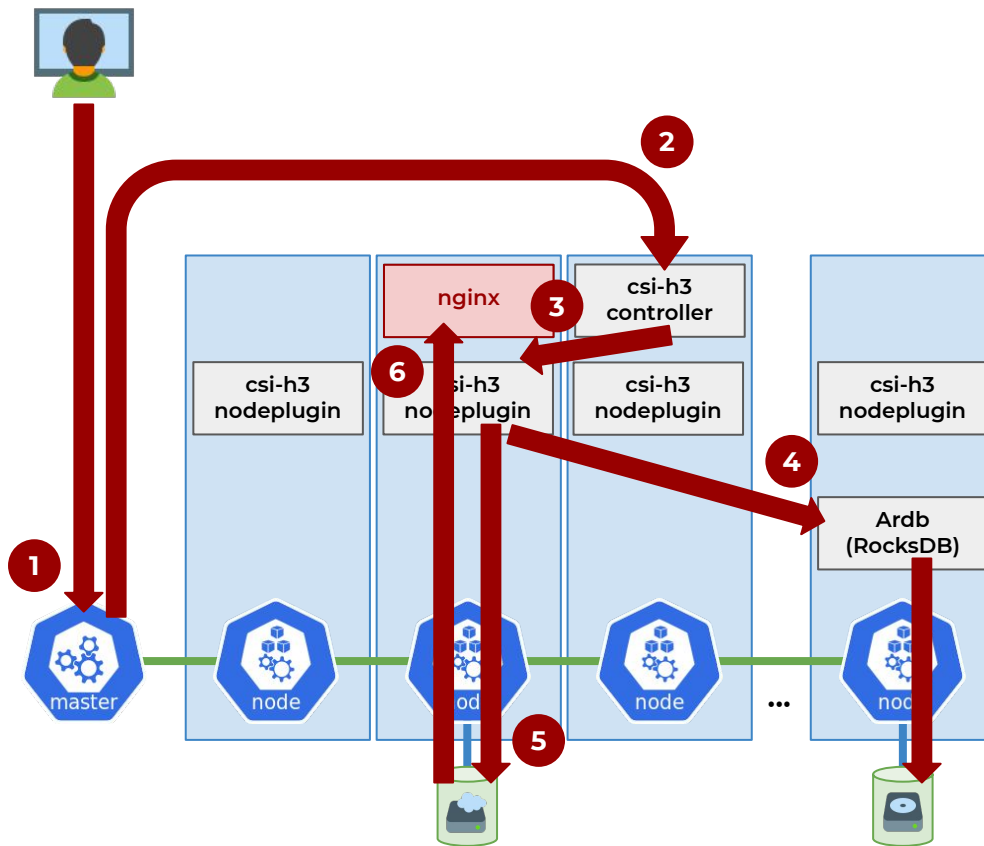
- 1 deploy service
- 2 need volume
- 3 create volume
- 4 create bucket and mount



- 1 deploy service
- 2 need volume
- 3 create volume
- 4 create bucket and mount
- 5 mount volume to host



- 1 deploy service
- 2 need volume
- 3 create volume
- 4 create bucket and mount
- 5 mount volume to host
- 6 attach volume to container



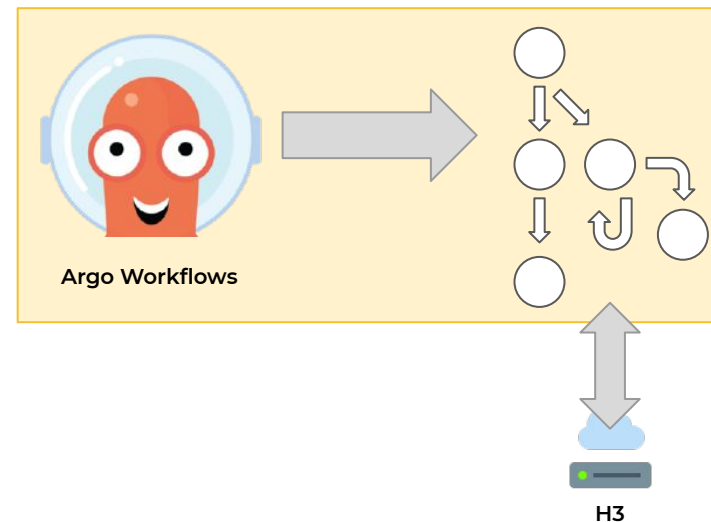
# Argo (s3-support branch)

Add support for H3 artifact repository

Get/put data interface using **h3cli**

Download artifacts from **Argo GUI**

Modified **Argo containers**



# Software releases



Description	Software	Images
H3 is an embedded object store in C, Python, and Java	<a href="https://github.com/CARV-ICS-FORTH/H3">https://github.com/CARV-ICS-FORTH/H3</a>	carvicsforth/h3
CSI driver for H3	<a href="https://github.com/CARV-ICS-FORTH/csi-h3">https://github.com/CARV-ICS-FORTH/csi-h3</a>	carvicsforth/csi-h3
Performance test for H3	<a href="https://github.com/CARV-ICS-FORTH/h3-benchmark">https://github.com/CARV-ICS-FORTH/h3-benchmark</a>	
Argo Workflows: Get stuff done with Kubernetes.	<a href="https://github.com/CARV-ICS-FORTH/argo">https://github.com/CARV-ICS-FORTH/argo</a> (clone of official Argo, look for h3-support branch)	carvicsforth/argocli carvicsforth/argoexec carvicsforth/workflow-controller



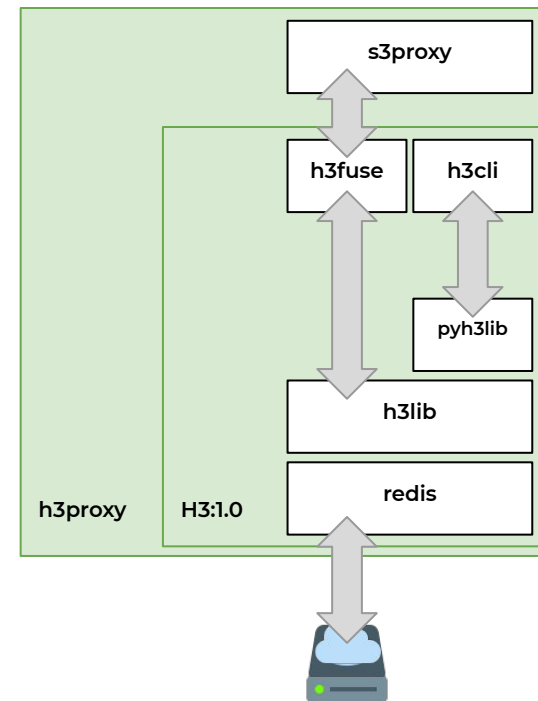
# Ongoing work

Even simpler Kubernetes integration with DLF\*

Expose S3 API (s3proxy to jClouds to H3)

Bucket and object **attributes**

**Compression** for data




\*<https://github.com/IBM/dataset-lifecycle-framework>

# H3 & csi-h3

Full ecosystem is open source\*

Actively developed

Funded as part of the EVOLVE H2020 project\*\*



\*<https://github.com/CARV-ICS-FORTH>

\*\*<https://evolve-h2020.eu>