Storage Quality-of-Service in Cloud-based Scientific Environments: a Standardization Approach

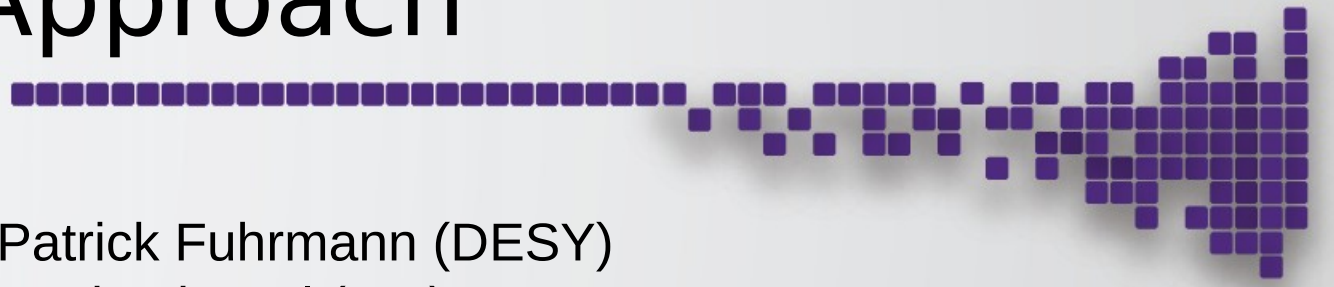

Patrick Fuhrmann (DESY)
Benjamin Ertl (KIT)
Maciej Brzezniak (PSNC)
Paul Millar (DESY),
Marcus Hardt (KIT),
Vladimir Sapunenko (INFN-CNAF)
Giacinto Donvito (INFN-Bari)
Andrea Ceccanti (INFN-Bari)

CHEP 2016 at San Francisco, USA
2016-10-11
https://indico.cern.ch/event/505613/contributions/2230920/

INDIGO-DataCloud is co-founded by the Horizon 2020Framework Programme

QoS: expectations and promises

INDIGO - DataCloud

Users

Storage behaves how I expect

Storage providers

Promises on how storage behaves

QoS: why bother?

“We have a finite budget”

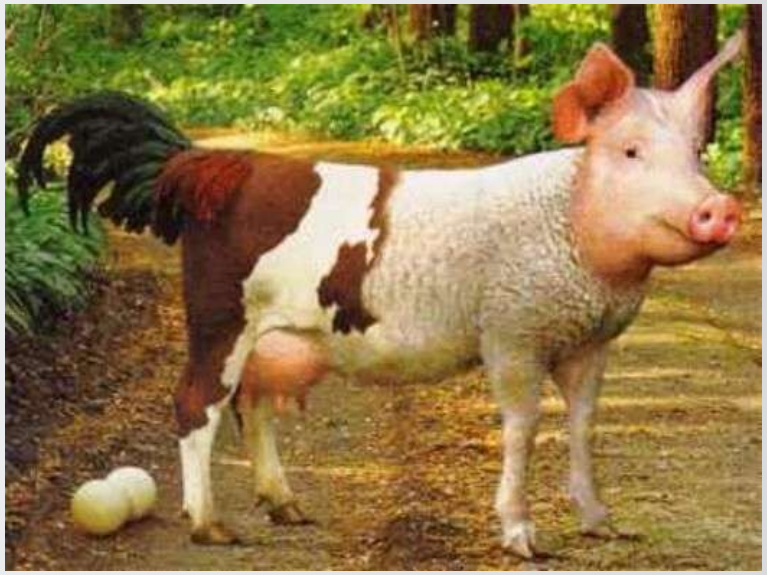

“None of the available storage meets all requirements”

“Chose any two” toy model

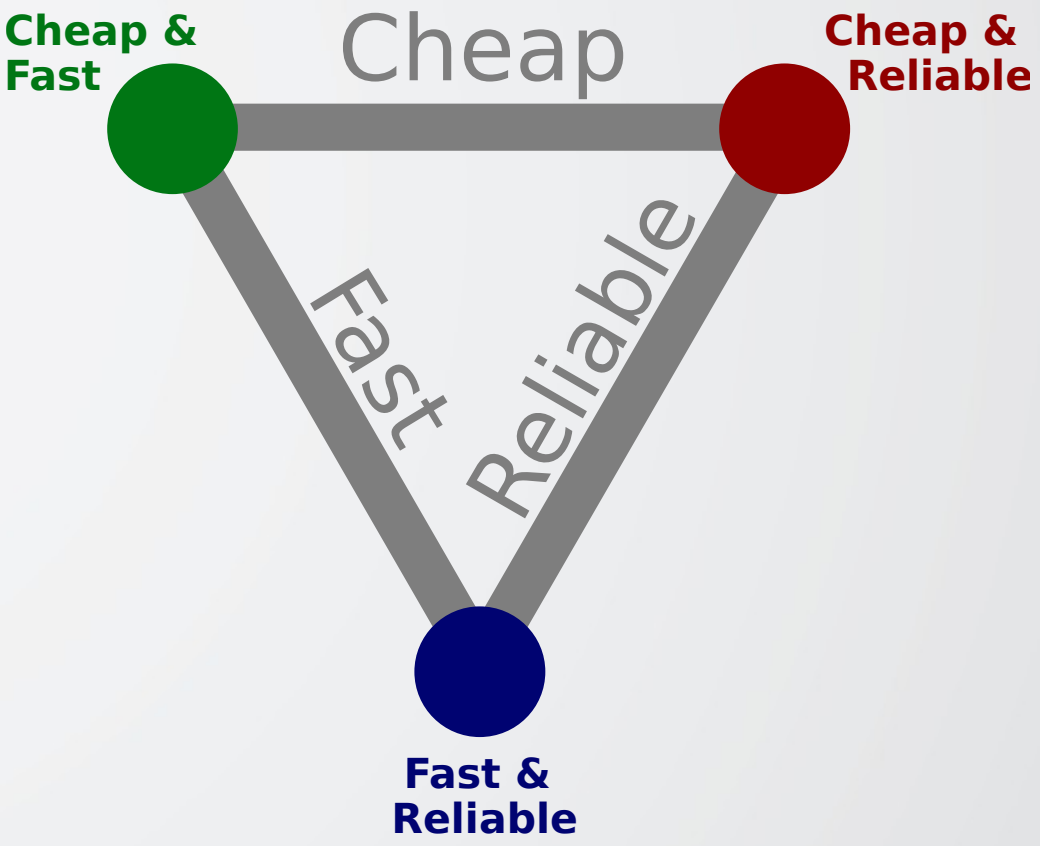

Existing QoS in HEP

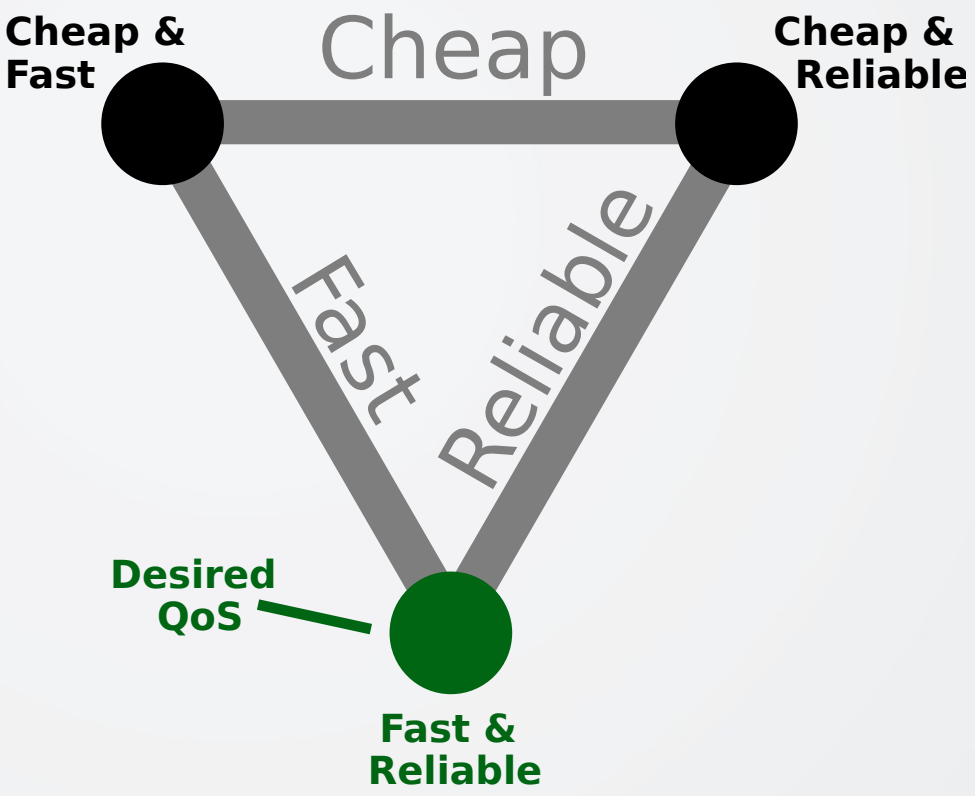

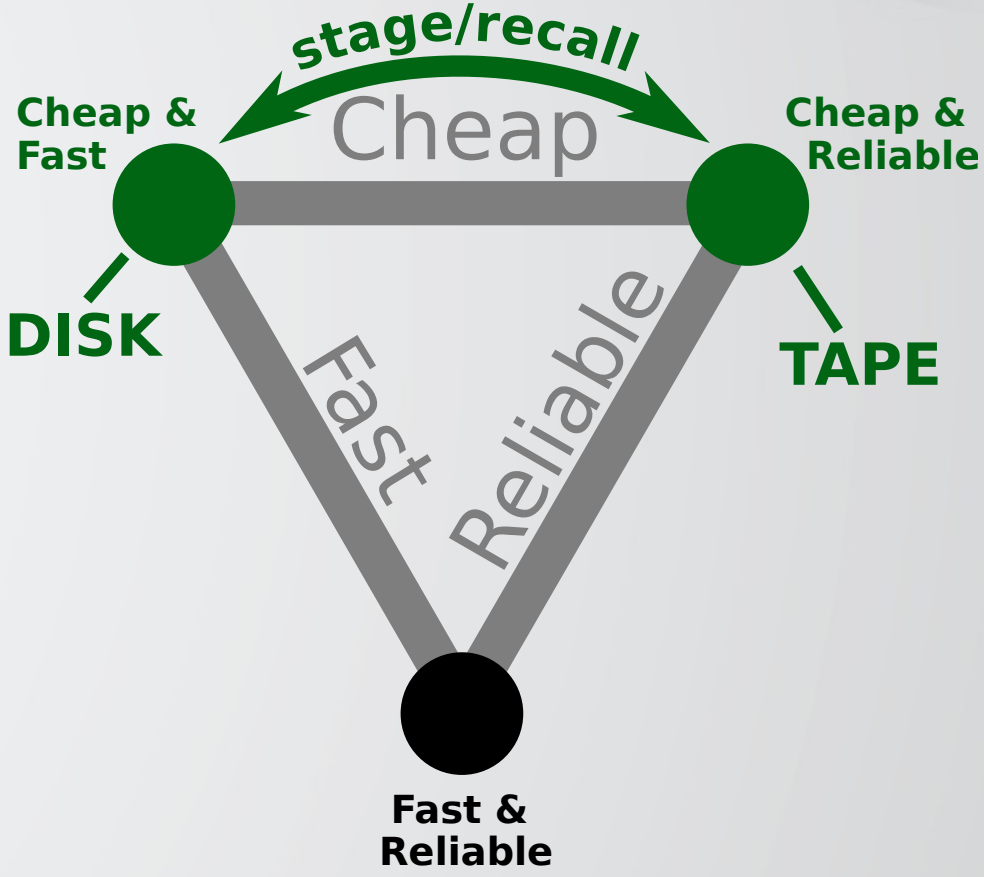

Possible directions

(Example set)

DISK

TAPE

- **SCRATCH**
- **FAST (the need for speed)**
- **OUTPUT (external redundant copies)**
- **OUTPUT (not yet redundant)**
- **LOW-COST (latency not an issue)**
- **ARCHIVAL (expensive to recreate)**

Why?

"We have a finite budget"

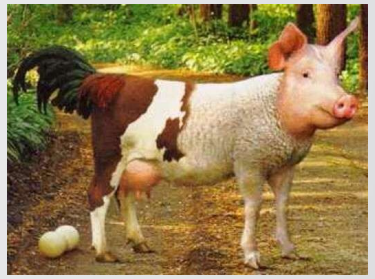

"None of the available storage meets all requirements"

Choosing QoS

INDIGO - DataCloud

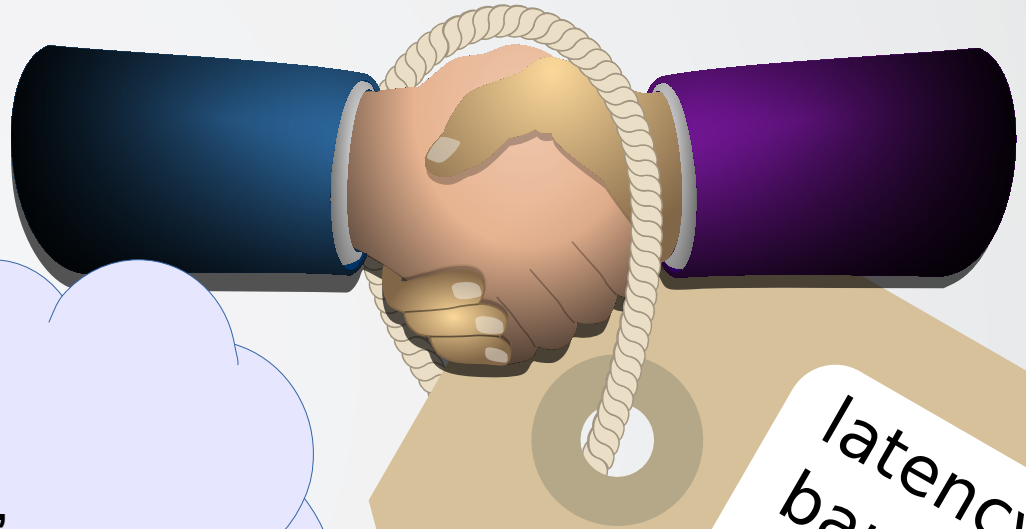

latency: ....
bandwidth: ...
durability: ....
cost-model: ...

"**SCRATCH**"
(latency)

"**ARCHIVAL**"
DURABILITY

"**FAST**"
LATENCY &
BANDWIDTH

User expectations

QoS #1: **SCRATCH**

QoS #2: **SCRACH, FAST**

QoS #3: **ARCHIVAL**

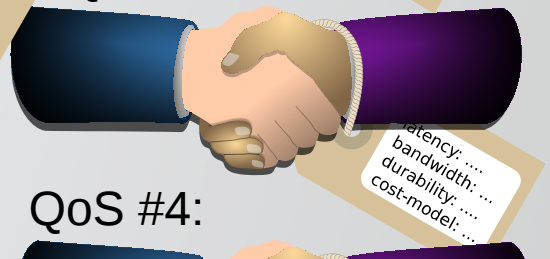

QoS #4:

Controlling QoS with CDMI

SNIA

Cloud Data Management Interface (CDMI™)

Version 1.1.1

ABSTRACT: This CDMI International Standard is intended for application developers who are implementing or using cloud storage. It documents how to access cloud storage and to manage the data stored there.

This document has been released and approved by the SNIA. The SNIA believes that the ideas, methodologies, and technologies described in this document accurately represent the SNIA goals and are appropriate for widespread distribution. Suggestion for revision should be directed to http://www.snia.org/feedback/.

With some extensions:

- **Modifying QoS (from SNIA)**
- **Discovering allowed QoS transitions**
- **Monitoring QoS transition progress**

CDMI Reference implementation

- **Query** available QoS options.
- **Control** QoS of files and directories

  Does not support ingress/egress
- Uses **plugin architecture** to support different backends:

  Support for dCache, S3/CEPH, HPSS, StoRM/TSM, GPFS.
- **Discover** alternative protocols

  for data ingress/egress

Defining QoS attributes

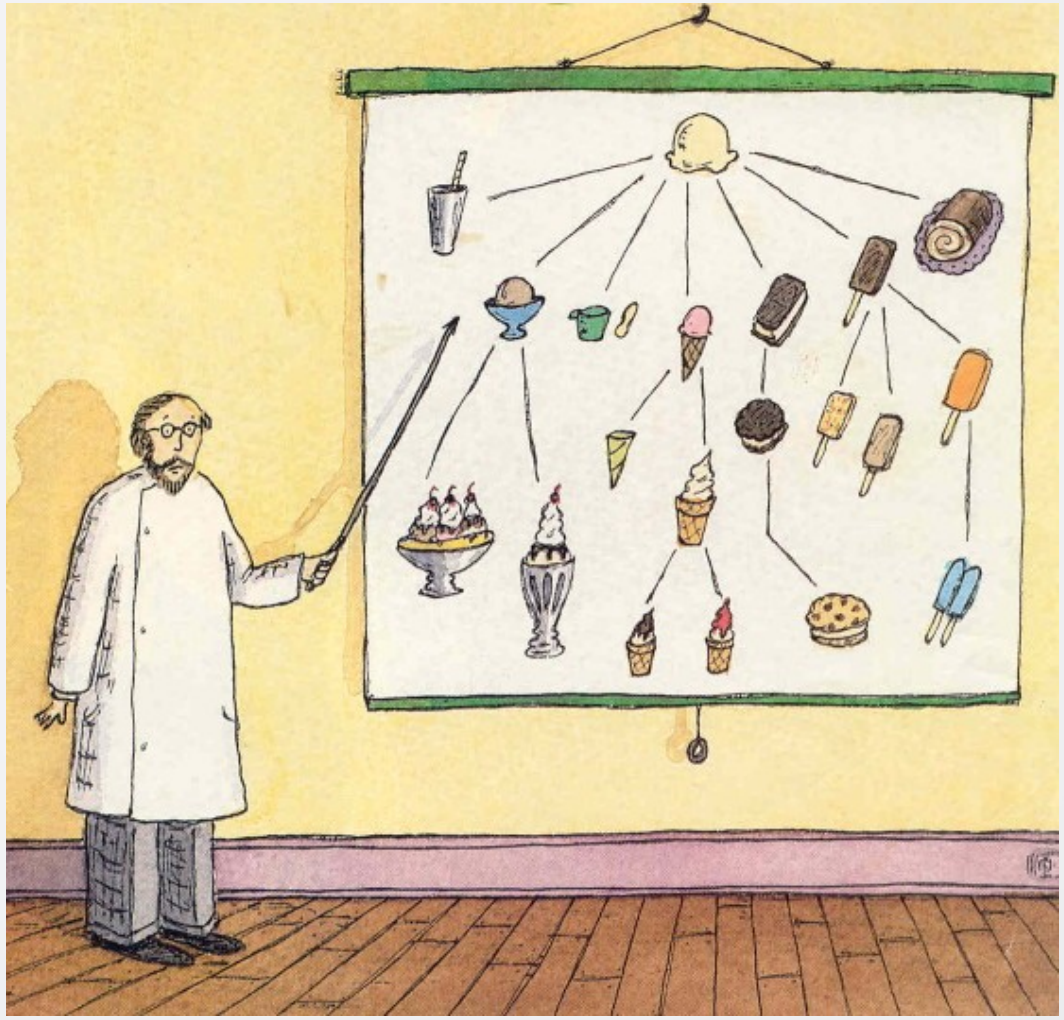

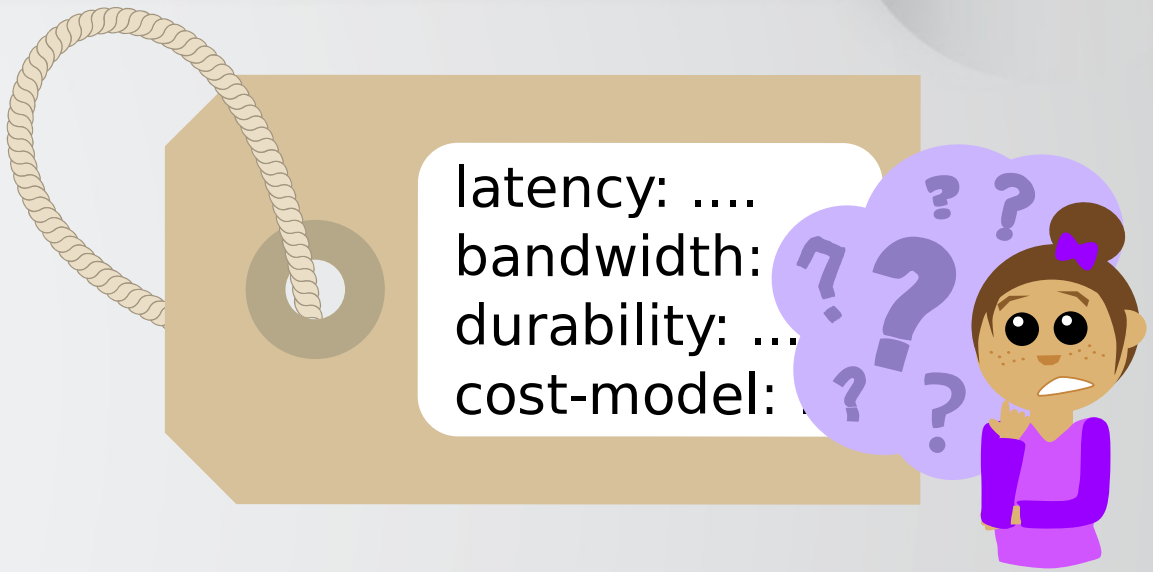

Next steps...

- Building up **test bed**

  Demo in November for EU review
- Complete **integration** with OneData.
- Add support for **custom attributes**.
- **Delegate authorisation** to INDIGO-DataCloud SLA.
- Roll out into **production**.

INDIGO - DataCloud

Backup slides

Toy model to more realistic view

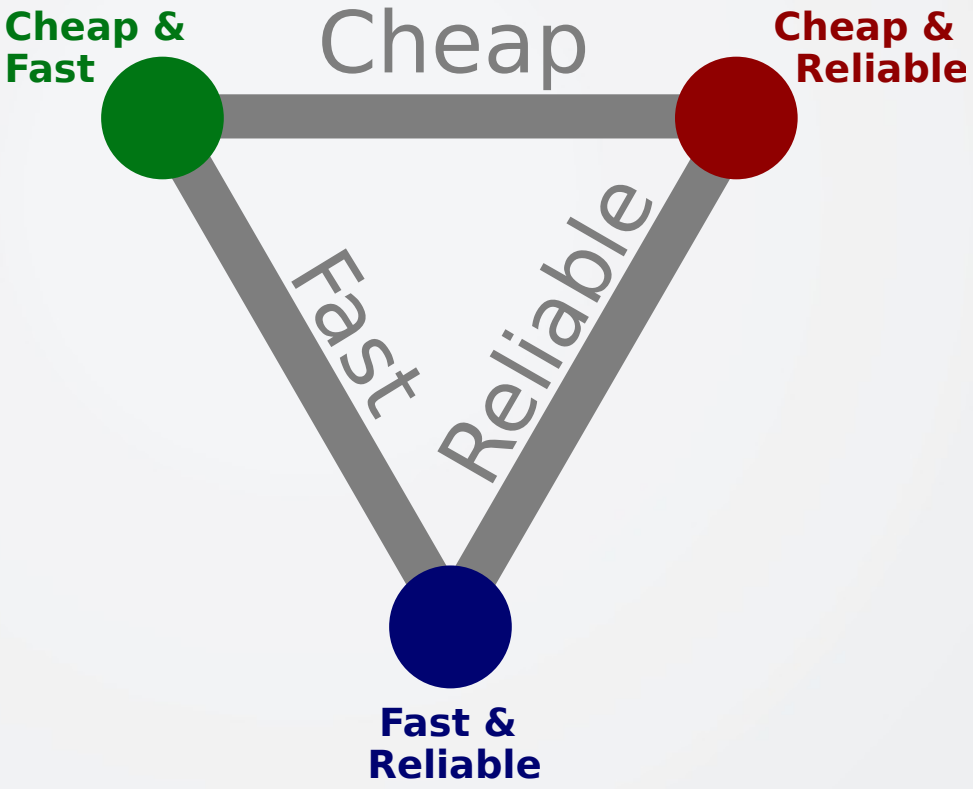

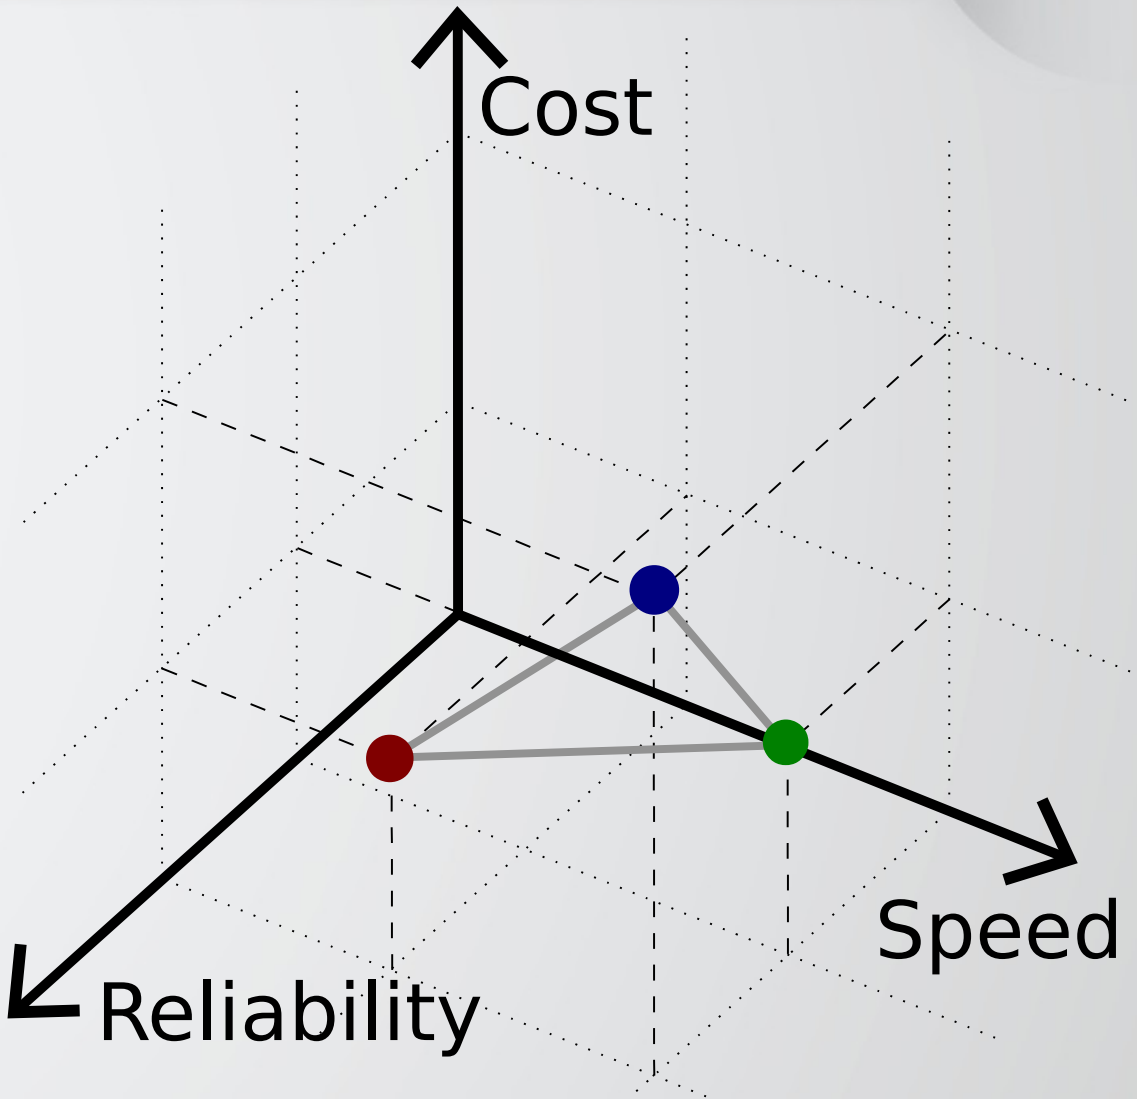

Example ARCHIVAL criteria