SGAS Accounting System Status and future plans

Josva Kleist
Software Coordinator, NDGF
EGEE'09, September 23rd, 2009

NDGF
Nordic DataGrid Facility

NDGF
Nordic DataGrid Facility

Overview

- SGAS status
- New requirements
-

NDGF
Nordic DataGrid Facility

Some History (of SGAS)

- Developed by Swedish Researchers
  - Around 2005-2006
  - Have all left academia
- Bank and Usage Tracking
  - Fused together in the same product
- Based on Globus WSRF
- ”Research prototype”

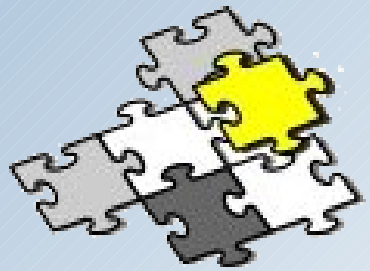

SGAS

NDGF
Nordic DataGrid Facility

Why use SGAS?

- It works with ARC (and GRAM)!
- Reasonably generic.
- Generates OGF UR (well sort of).
- SGAS is deployed and used by NDGF
  - Also used by M-Grid and SweGrid.

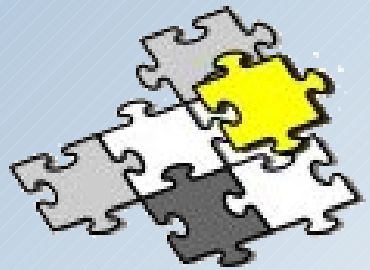

SGAS

NDGF
Nordic DataGrid Facility

Existing setup

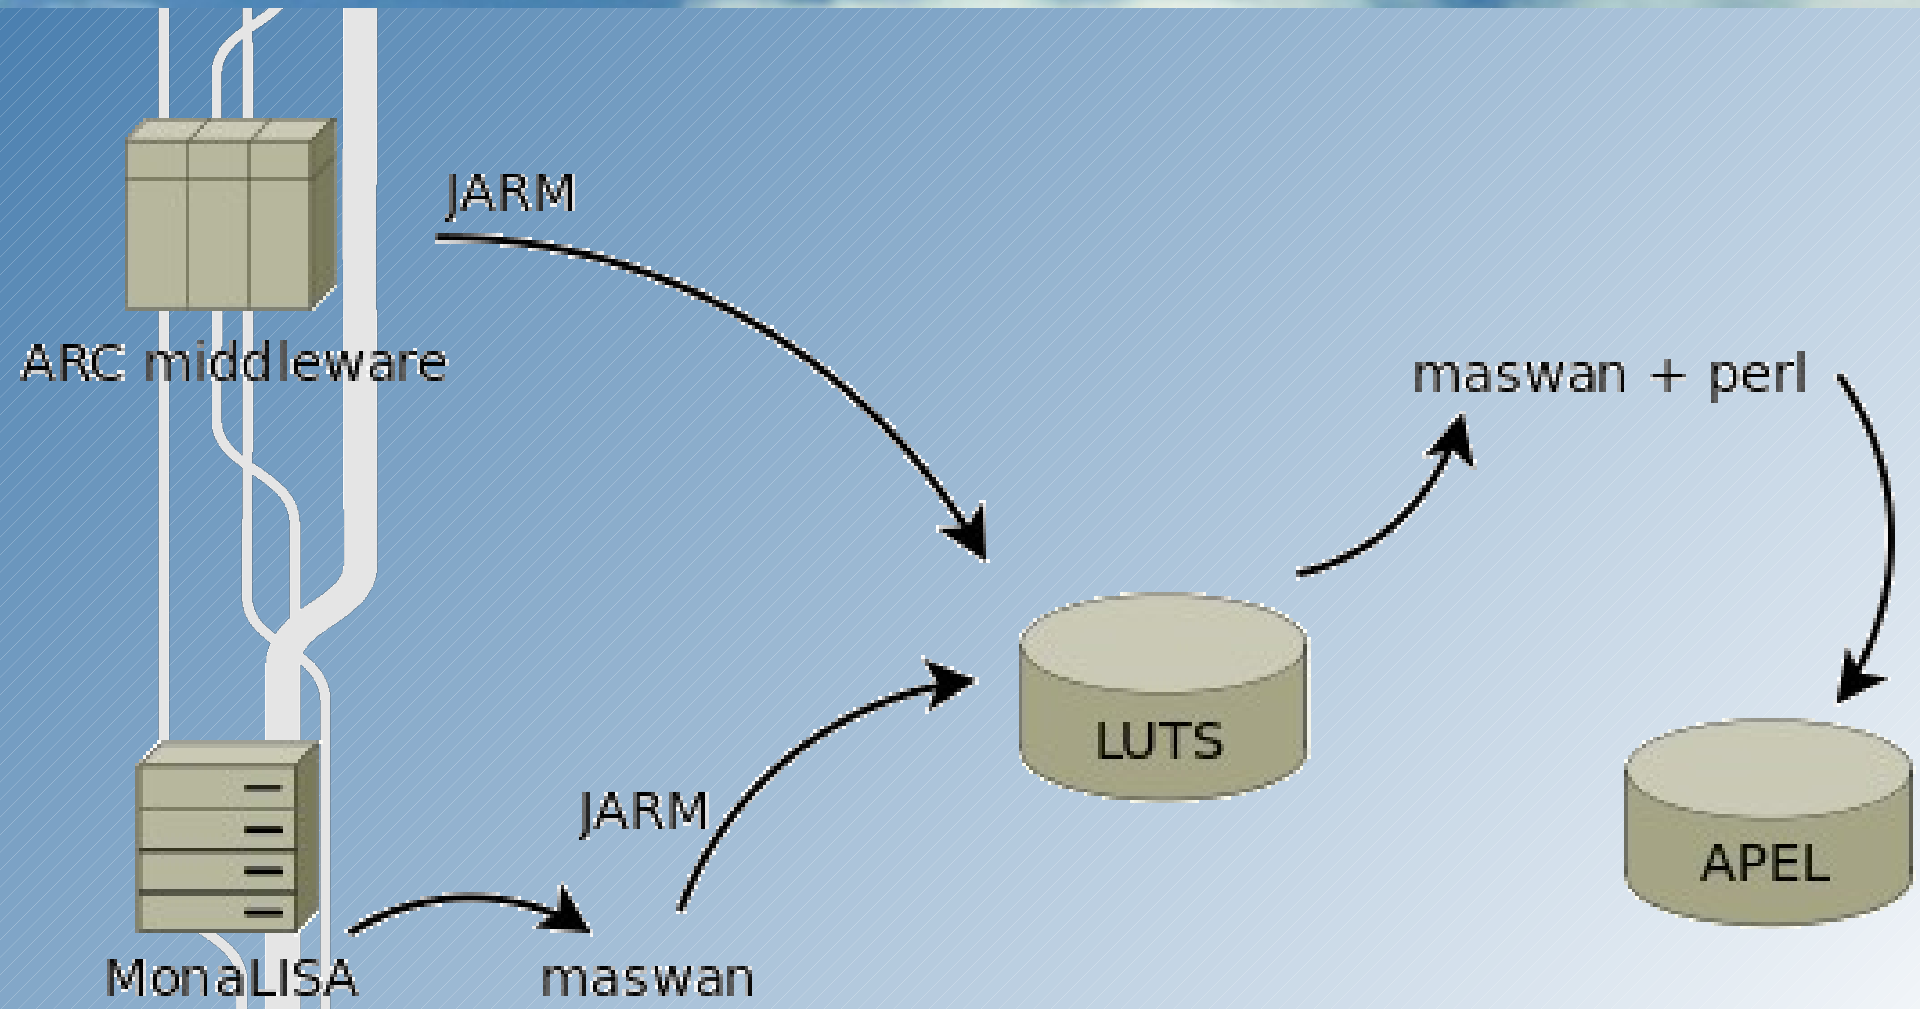

Problems

- No one really knew how it worked.
- Embedded XML database (gahh...)
  - Data extraction is difficult / silly.
- Heavy use of Globus WSRF stack
  - Big and unwieldy, tricky to understand.
- Only claiming to use standards.
- No active development.

NDGF
Nordic DataGrid Facility

Globus WSRF Stack

Description

Actual functionality

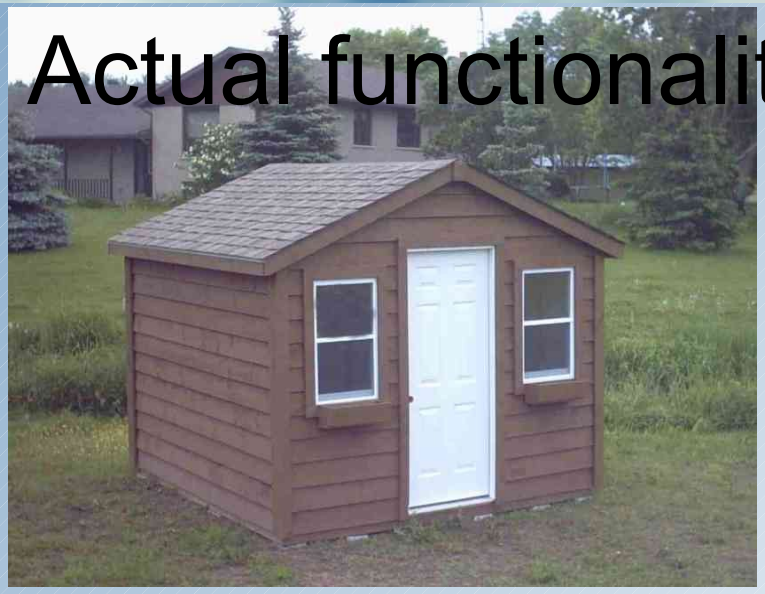

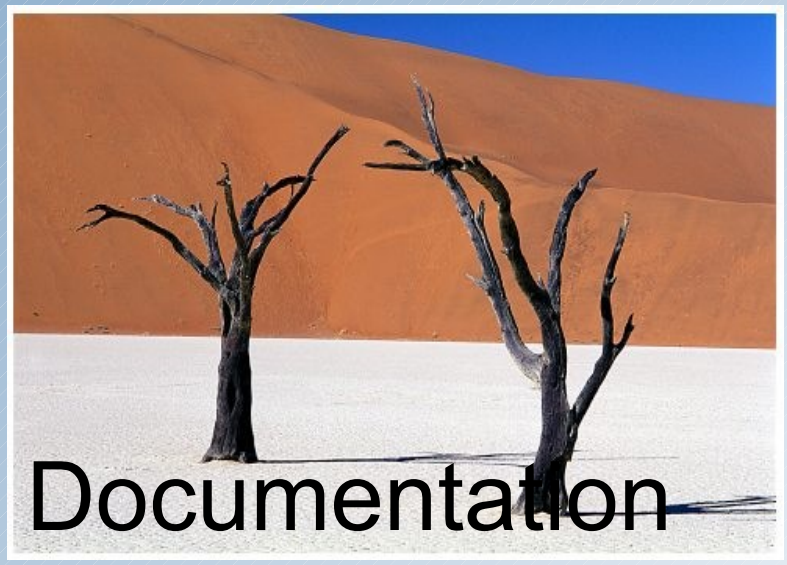

Documentation

NDGF
Nordic DataGrid Facility

New Requirements

- Remove manual steps for reporting and interfacing.
- Allow federated setup:
  - Report data in several places,
  - Aggregate data from several places.
- Use real OGF UR (whatever that means).
- Simplify deployment at sites.
- Ease development:
  - Get rid of Globus.

NDGF
Nordic DataGrid Facility

Envisioned setup

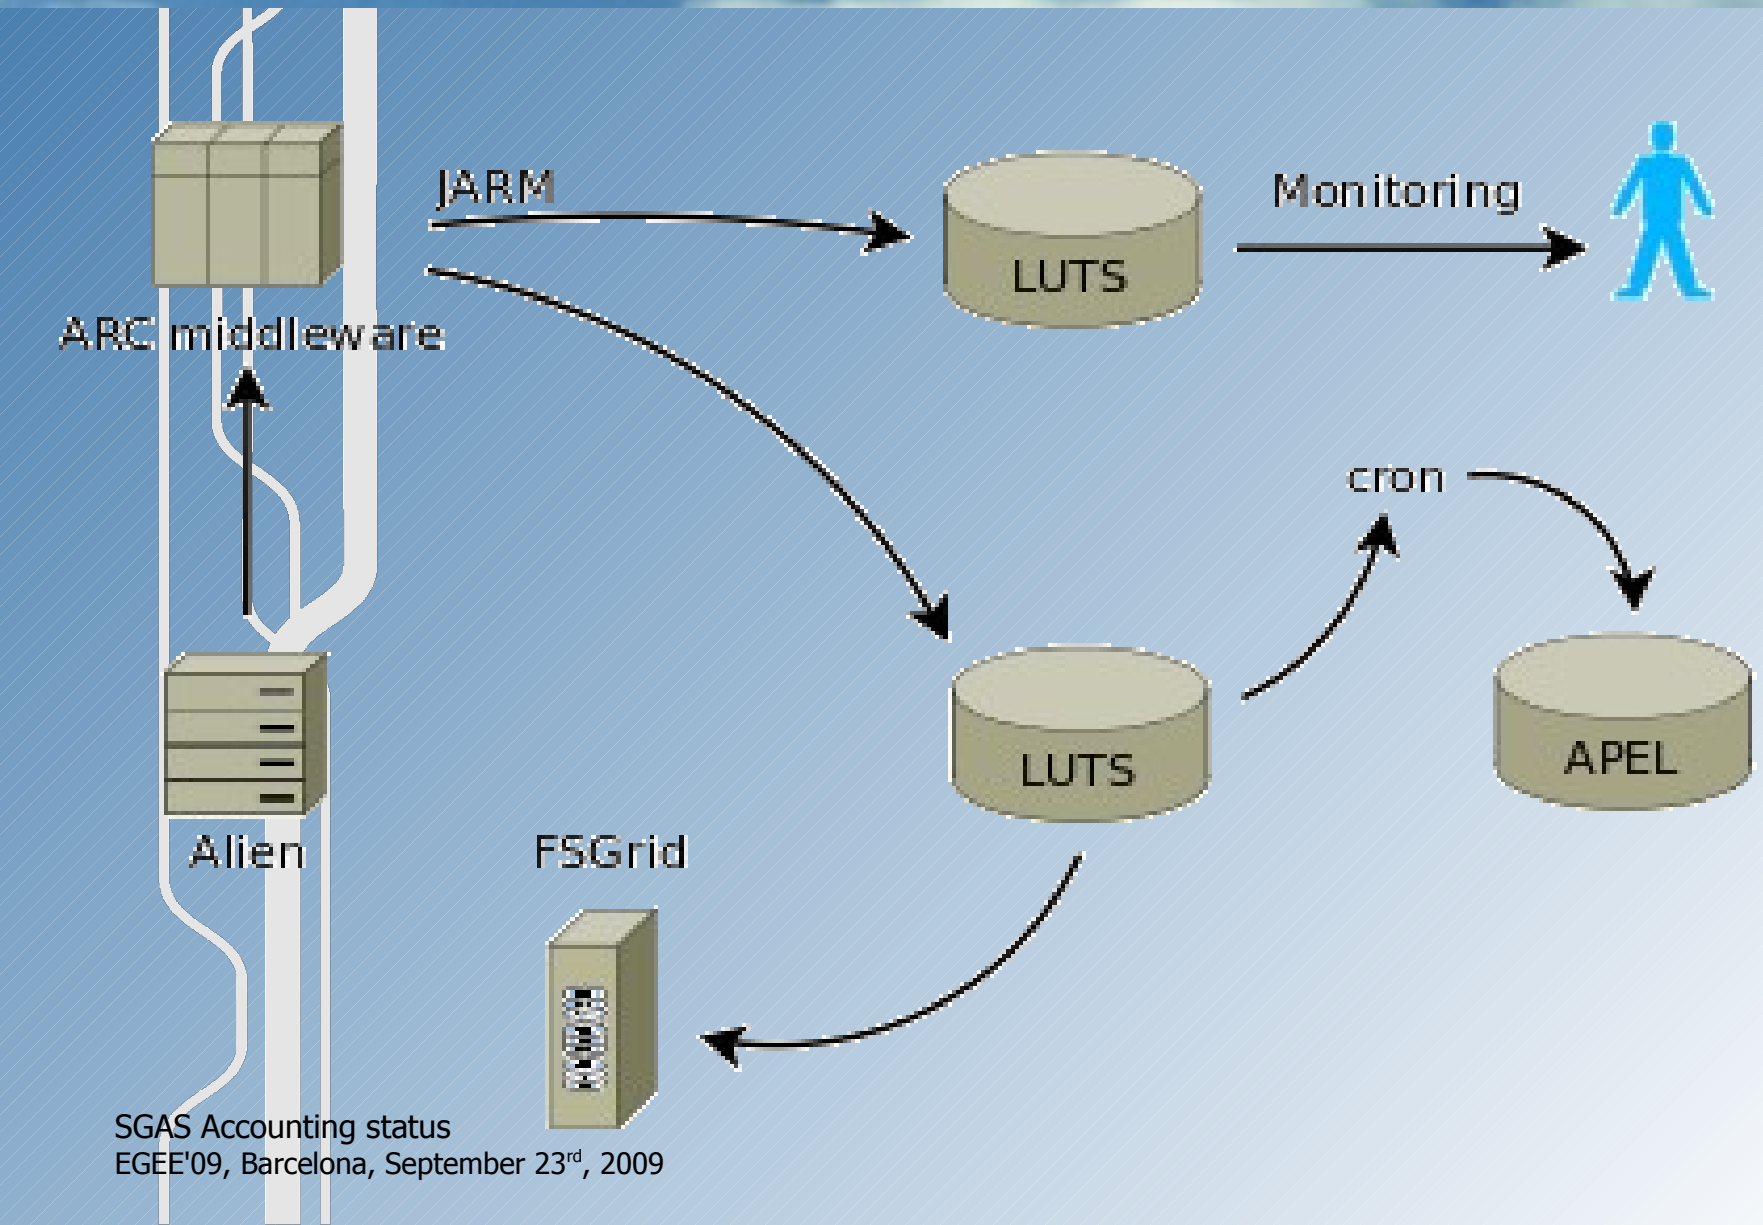

NDGF
Nordic DataGrid Facility

Why more than one LUTS?

- Legislation – some information are not allowed to be exported from a country unless user has given consent
- Store data where it is relevant:
  - Sites stores data about jobs run at site
  - VO stores data about jobs run by their members
  - Country stores data about jobs run in the country, or by its residents.
- Can be kept in one database.

NDGF
Nordic DataGrid Facility

SGAS Development

- Restarted in December
- One NDGF developer
  - Gird.se, grid.dk, KnowARC
- Development
  - New build system
  - Bank-LUTS split
  - Have SGAS use cron
  - New back-end database

NDGF
Nordic DataGrid Facility

Rewriting SGAS

- Two security holes where found before summer:
  - 1. Bad protocol design (and Globus)
  - 2. Insane Globus default (fixable).
- Core SGAS functionality rather small.
- An almost-from-scratch implementation was started in August:
  - Two weeks to re-create original functionality.

NDGF
Nordic DataGrid Facility

The Database tale

- Originally eXist was used embedded in SGAS:
  - Custom locking layer in SGAS
  - Performance was bad, working with data worse.
- Sedna:
  - External XML database, concurrent access supported,
  - Flaky stability, performance still bad.
- CouchDB:
  - External JSON database (schema less),
  - Data conversion required, performance is good.

NDGF
Nordic DataGrid Facility

Current Status

- Registration pipeline is working
  - Two servers + three loggers running.
- External database – yay!
- Experimental views.
- Current activities:
  - VO information in usage records,
  - WLCG reporting details (per user reporting).

NDGF
NORDIC DATAGRID FACILITY

SGAS View

View - Shiretoko

File Edit View History Bookmarks Tools Help

https://orval.grid.aau.dk:6143/sgas/view/wdmn

activemq

20.13. xml.etree.Eleme... | Twisted Documentation:... | Introduction to CouchDB... | View | View

Accumulated wall duration (hours) / Day / Machine Name

	2009-08-20	2009-08-21	2009-08-22	2009-08-23	2009-08-24	2009-08-25	2009-08-26	2009-08-27	2009-08-28
csc	76.1	85.5	157.4	90.3	134.8	166.9	52.3	109.2	29.3
alice-vobox.phys.jyu.fi	2.8	21.9	25.2	23.2	33.1	20.6	12.3	15.8	4.3
alice.lunarc.lu.se	31.5	281.8	319.0	401.1	532.4	383.0	89.8	300.7	94.1
alien-ce.smokerings.nsc.liu.se	43.9	240.7	361.7	452.9	232.3	400.8	211.1	207.5	43.6
benedict.grid.aau.dk	1.5	35.1	42.6	51.9	35.6	38.8	17.7	26.8	
francis.grid.aau.dk			16.0	2.1	9.5	8.6	12.1	11.3	0.4
grid.bccs.uib.no	35.0	213.8	259.8	248.4	226.2	250.7	188.8	199.2	76.8

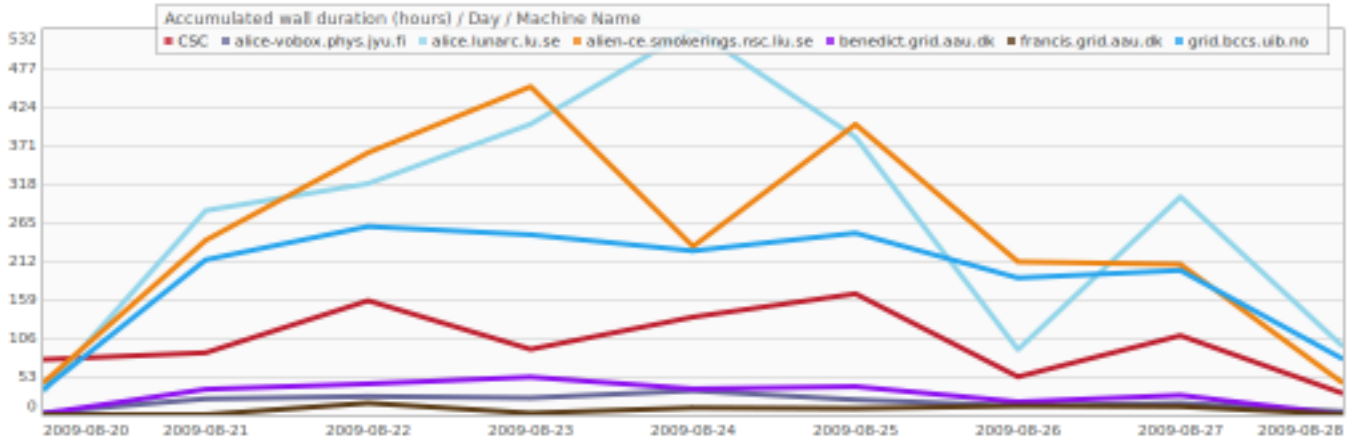

Done

NDGF
Nordic DataGrid Facility

Middleware Issues

- ARC:
  - Getting the right number of cores,
  - Getting start/end time from the LRMS,
  - LRMS scripts needs to be changed,
  - Getting VO information is tricky.
- Alien:
  - Everything is run under the same user (the pilot job problem),
  - Wish to register per-user.

NDGF
Nordic DataGrid Facility

Future

- Short-term
  - Deployment of loggers and server
  - Reporting to APEL
  - LHC starts in November
- Long-term
  - Monitoring graphs / tools
  - Eliminate silly work flow
  - Grid-wide fairshare

NDGF
Nordic DataGrid Facility

And more long term

- Same accounting system for all computational resource use at sites.
  - Local use through SSH
  - Grid use
  - Web access to services
- Tied up to federated identity providers.

FSGrid (bonus slide)

- Bring fairshare to the grid
  - Clusters have had fairshare for some time
  - Knowing resource-consumption is a prerequisite
- Model developed by Swedish Researchers :-)
  - Model is sound and fairly simple
- Status: Worked on a laptop 3 years ago
  - Likely to require a lot of work

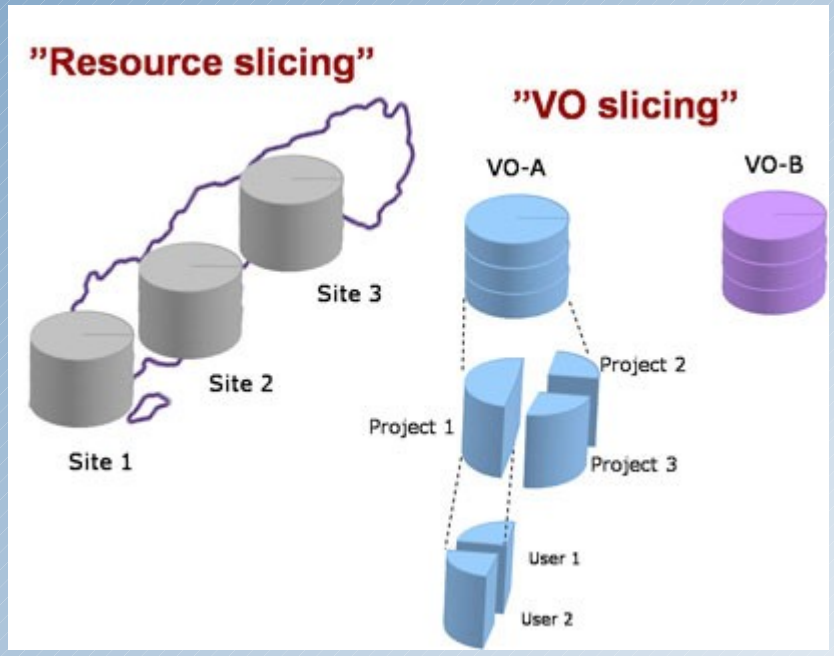

Questions

NDGF
Nordic DataGrid Facility