Grid computing for the SNO+ experiment

11/01/2013

Matthew Mottram

University of Sussex

Introduction to SNO+ & SNOLAB

Multi-purpose physics (neutrino) experiment at SNOLAB

SNO+ collaboration map

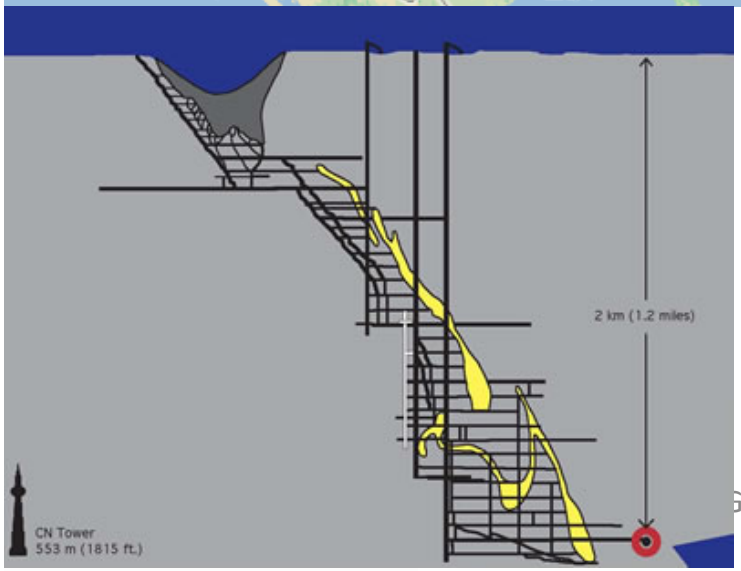

6800 feet (2 km) underground

Class-2000 cleanroom (in a working mine)

The SNO+ experiment

- Physics goals:
- **Double beta-decay**
- Solar, geo, reactor, supernova neutrinos
- Forbidden nucleon decay

2km underground (!!!)

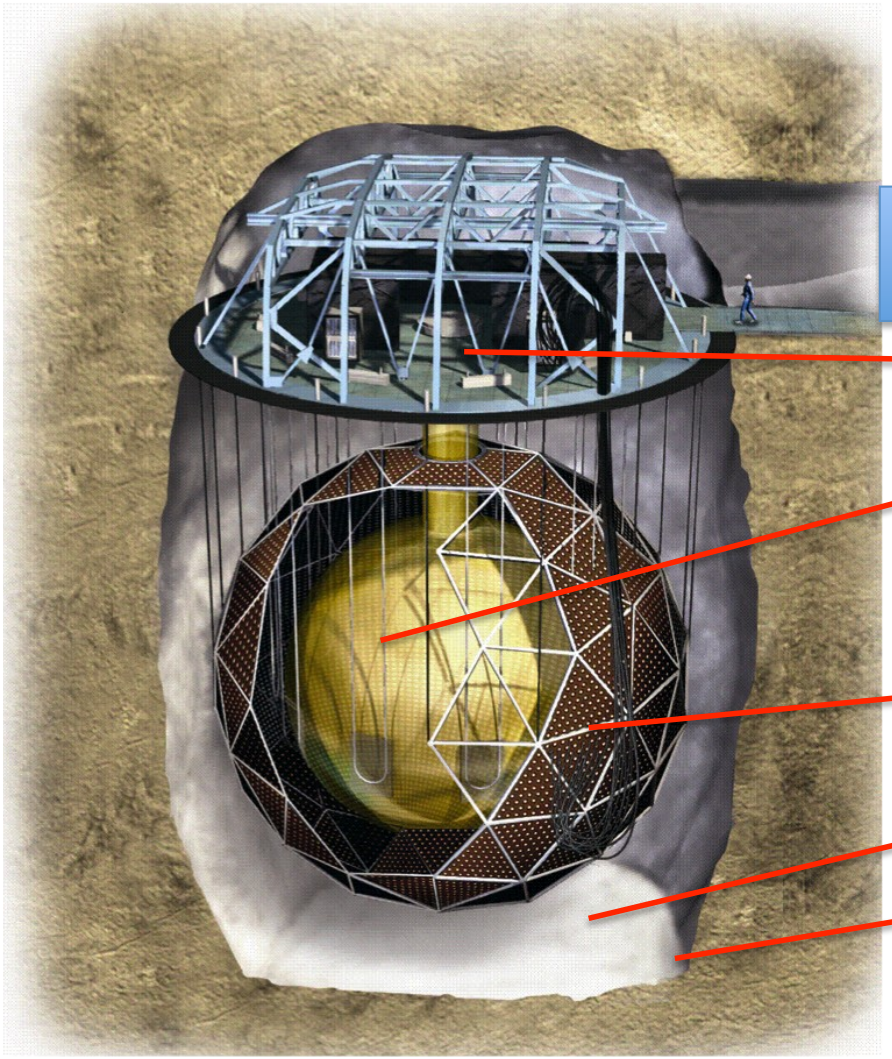

University of Sussex

SNO+

Data transfer & storage requirements

- Max data taking rate of 450Mbit/s (!)
  - Expect 10s TB of real data a year
  - Require 10x simulated data
    - Larger data-structure size compare to real data
    - But not all real data requires simulating
    - -> 10s-100 TB a year
- Very limited computing resources at SNOLAB:
  - Buffer data underground and on surface
  - Ship all real data off-site
  - All simulations processed off-site
- **Expect** (guess) **we need 10s of cores continuous use and 10s-100s of TB of storage a year to process/simulate required datasets.**

Resource layout

- SNO+ will store and process/simulate data at two locations:
- WestGrid (ComputeCanada):
  - Canadian Grid Tier 1 storage resource
  - One single batch farm (qsub submission)
  - 100 nodes, priority access
- LHC Computing Grid (LCG):
  - UK Grid Tier 1 storage at RAL
  - Processing at distributed Tier 2 sites
  - 1000s of nodes, but low priority
  - Use glite software and workload management services (WMS)
    - Do not submit jobs directly to specific site
    - Place requirements on data, software etc, WMS does the rest

LCG data-processing

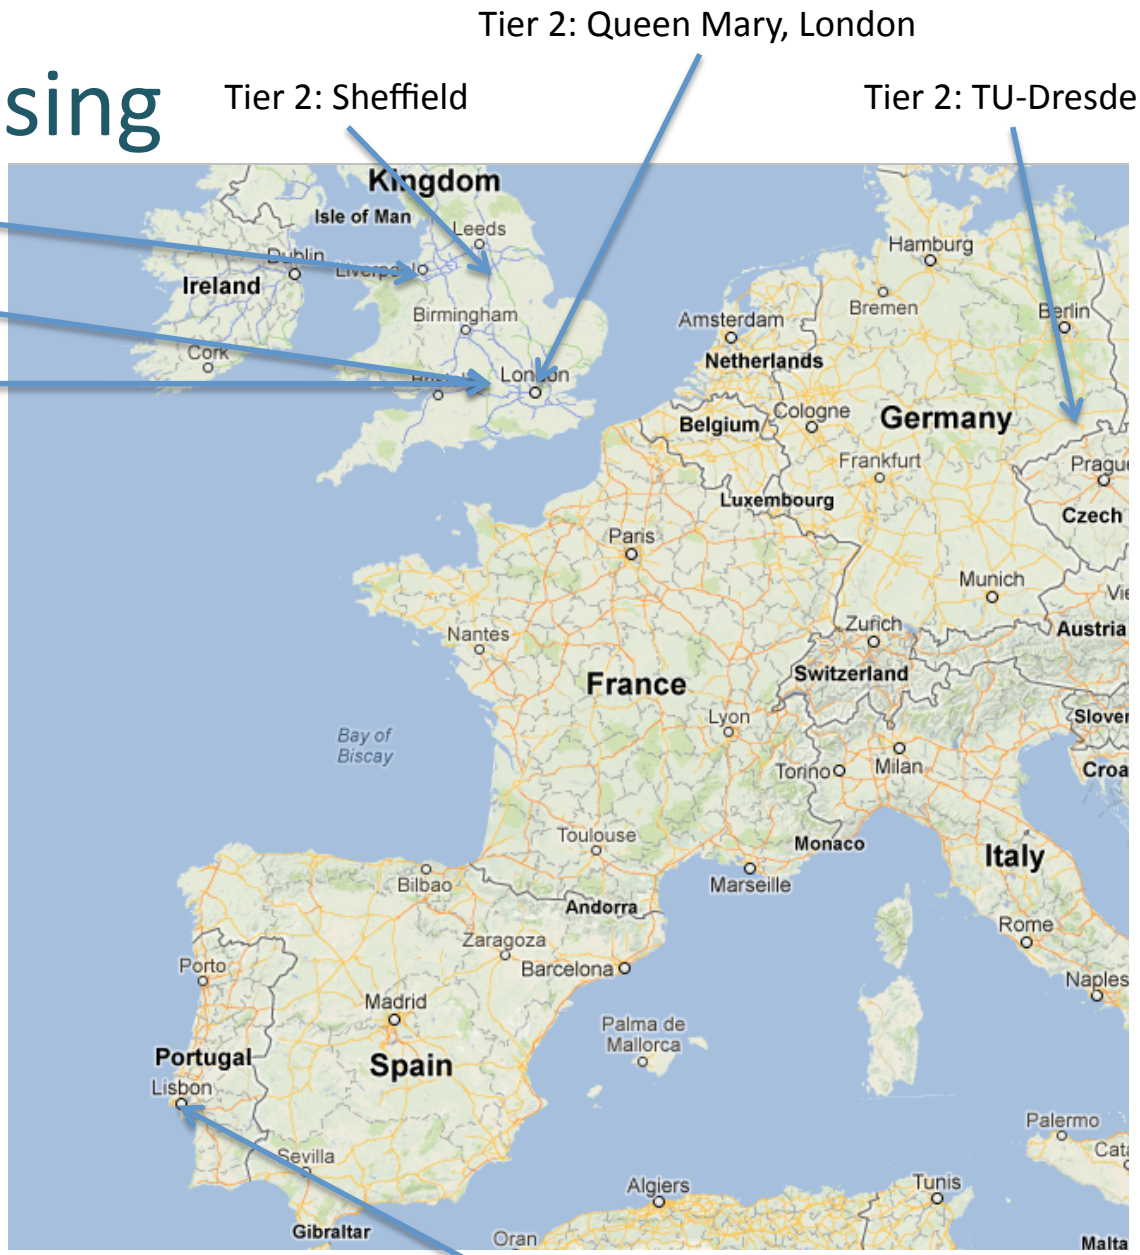

- Access to 7 sites
- Process/produce data at Tier-2 sites
- Archive data at Tier-1 storage
- Ship data out to Tier-2 sites ready for (re)-processing

University of Sussex

Grid job lifecycle

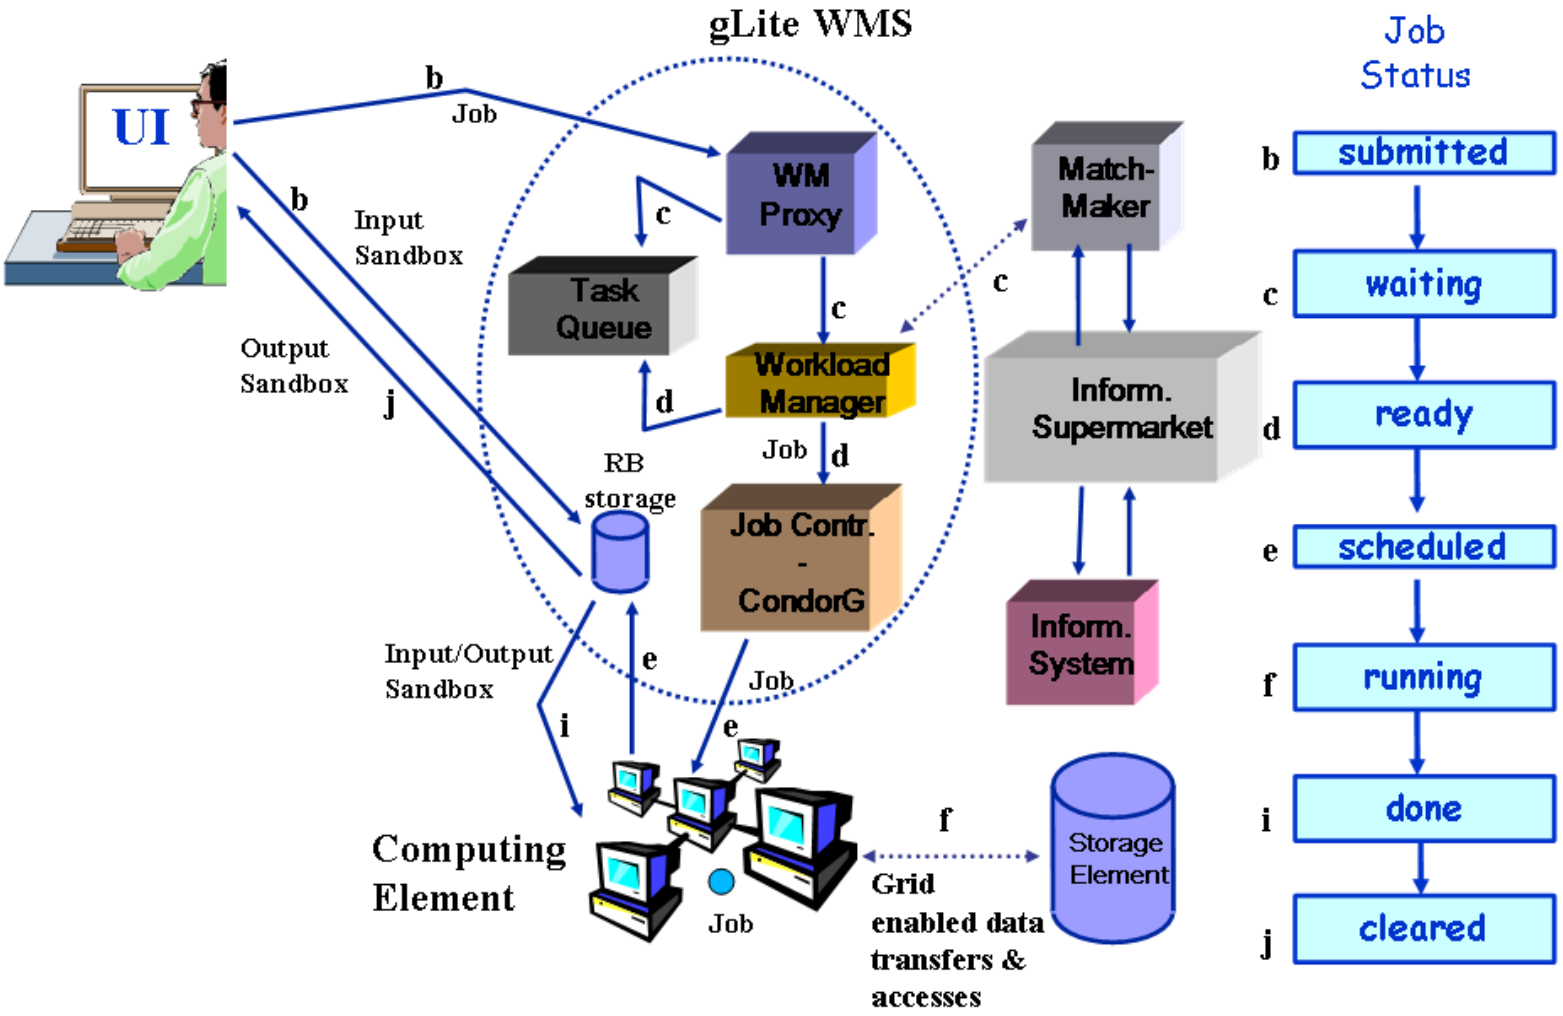

From glite-3.2 user guide

Grid jobs/data – current SNO+ method

- Submit job to Workload Management System (WMS)
  - WMS matches job requirements to available computing elements (CE) - software, memory etc
  - Submits job to specific queue
  - SNO+ rely entirely on the WMS for job management (so we use three)
- Jobs output data to their local storage element (SE)
  - All data logged in LCG File Catalogue (LFC)
  - SNO+ jobs *would* rely entirely on the LFC
    - But access only to one of these, at RAL
    - Official Monte Carlo and processed/raw real data are logged in a SNO+ database

Processing & Production

- Two very different production environments
  - Plus possibility of introduction of more farms/grids
  - Plus smaller analysis groups may want to reprocess data at local institution
- -> Need for backend agnostic production system

- Ganga: “configure once, run everywhere”
  - Python job management software
  - Setup to run SNO+ software on range (batch, grid etc) of submission backends, low user expertise required

Ganga vs no Ganga

Ganga

ganga_script.py

```
j=Job()
j.application='RATUser'
j.application.ratMacro='mymacro.mac'
j.application.ratBaseVersion='4.1'
j.application.outputDir='myDirectory'
j.backend='LCG'
j.submit()
```

job_description_file.jdl

```
Executable="job_script.scr",
Arguments={"-l","rat.log"},
StdOutput = "job.out",
StdError = "job.err",
InputSandbox = {"job_script.scr"},
OutputSandbox = {"rat.log","job.out","job.err"},
VirtualOrganisation = "snoplus.snolab.ca"
```

job_script.scr

No Ganga

```
source $VO_SNOPLUS_SNOLAB_CA_SW_DIR/snoing-
install/env_rat-4.1.sh
rat mymacro.rat
lcg-cr –vo snoplus.snolab.ca –d
$VO_SNOPLUS_SNOLAB_CA_DEFAULT_SE –P myDirectory/
myoutput.root –l lfn:/grid/snoplus.snolab.ca/
myDirectory/myoutput.root myoutput.root
```

Ganga vs no Ganga (2)

Job submission and monitoring

```
j.submit()
j.status()
```

Job submission and monitoring

```
glite-wms-job-submit –c config.conf –a –o myjob.jid
job_description_file.jdl
glite-wms-job-status –i myjob.jid
glite-wms-job-output –i mujob.jid
```

- GangaSNO plugin created:
  - Jobs on batch/grid/local node will automatically setup the correct environment to run SNO+ software
  - User just provides a job (Geant4) macro
  - Very little configuration required by user, but ability to configure is there if desired.

Production systems

- SNO+ will automate processing and simulation of "official" SNO+ data:
  - Production systems with a big bookkeeping task:
    - O(10k) data runs per year (real data)
    - Number of processing modules (~10) per run
    - One job per module
    - + 10x Monte Carlo with 10s of modules per physics run
  - Use CouchDB database:
    - Schema free
    - Documents, not tables
    - Documents stored as JSON objects
    - Pre-processed ("design") views – fast querying

Production systems

Concept

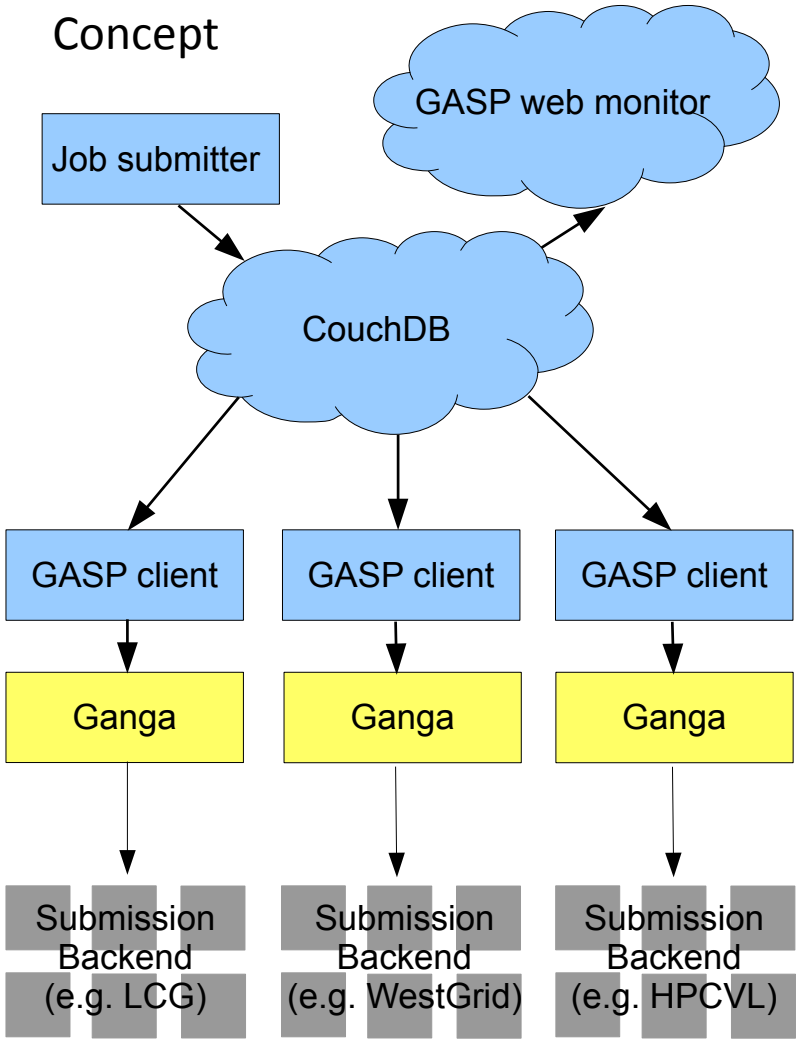

Database (CouchDB) document


```
{
 "_id": "19bfec64083952f4f6c2bd37df00110f",
 "_rev": "1-7b833b61ceccc06acedf9be72eabcb51",
 "status": "waiting",
 "run": 13,
 "module": "Dummy",
 "locked": null,
 "type": "job",
 "passes": [
        { "status": "waiting",
          "ratv": "3_00",
          "subRun": [
                { "status": "waiting",
                  "failCount": 0,
                  "output":
                        { "cksum": null,
                          "loc": null,
                          "name": null,
                          "se": null,
                          "size": null },
                 "subRun": 0 },
         "pass": 0 }
        ]
}
```

SNO+ roadmap

- Preparing for data-taking:
  - Spring 2013: fill with water, first physics data
  - Autumn/Winter 2013: fill with liquid scintillator
  - 2014: deploy double-beta-decay isotope
- In the meantime (a small selection of recent/ongoing work):

Hold down rope net (AV buoyant)

Installation of in-situ calibration systems

Empty detector (air filled) commissioning

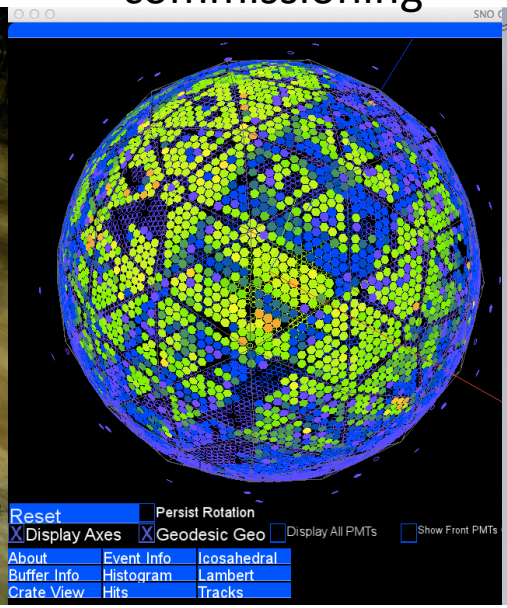

AV acrylic cleaning

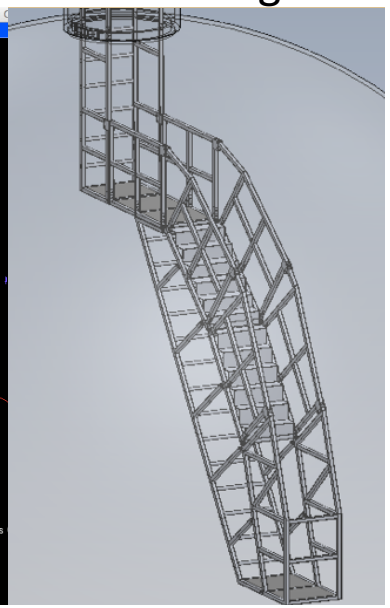

Thanks for listening!