

Experiment Input

From the experiments! Mostly.
Sometimes. Occasionally. Maybe.

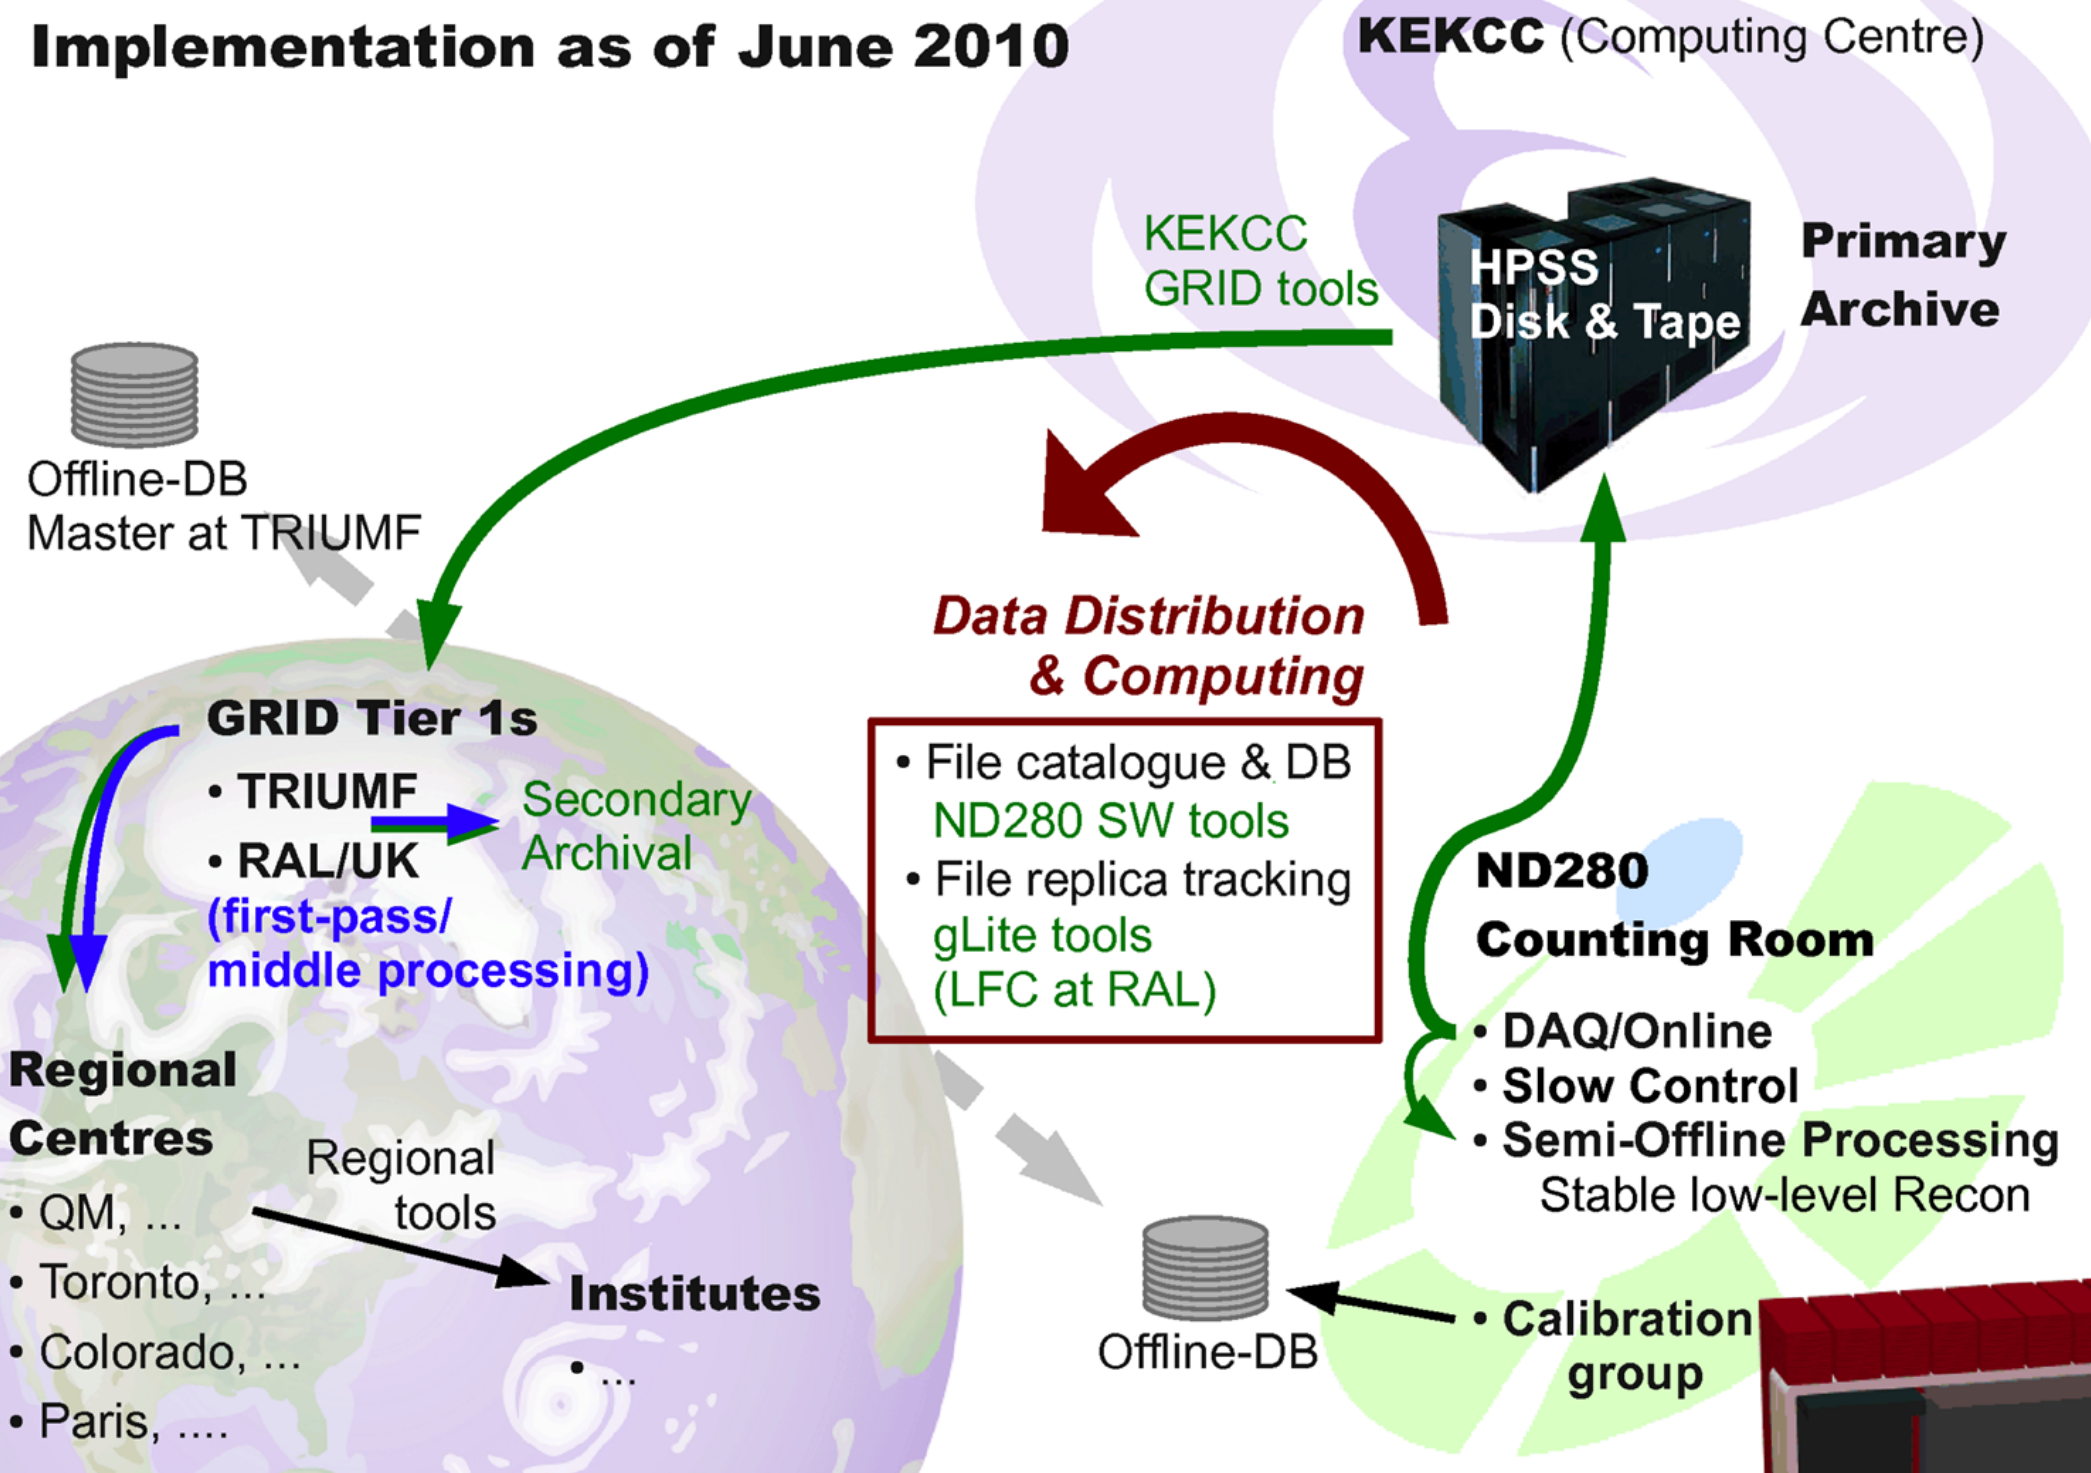
Overview

- CAMONT - Matt Doidge
- T2K - Chris Walker
- SuperB - Chris again
- MICE - Sam and Jens
- UKQCD - Wahid
- ATLAS - actually agreed to do some slides :-)
- LHCb - "nothing to report"
- CMS - didn't respond :-)

T2K

Chris Walker
on behalf of Ben Still

ND280 Data Distribution Implementation as of June 2010



UK Data (order of size)

- RAL

- Primary store 100% of data 20TB +

- QMUL

- Old primary store 8TB replication of RAL data

- Liverpool

- First Monte Carlo Production Stored here #TB,

- Lancaster

- Test beam data + some Monte Carlo #

- Imperial College

- Some duplication of processed files

- Other T2K institutes not utilised for storage yet

Questions

- FTS

- Automatic registration in LFC?

- Automatic distribution/replication of data from RAL?

- I am unsure how to get a job sent to the data

- Are there such things as data tags?

- How to I tell the optimum way to access data at different sites – turls, rfio etc and where is the information how to do this for each protocol?

SuperB

Chris, from Adrian Bevan

SuperB Grid Experience

- Setup: Core of ex CDF/CMS/... people set up/adapted existing tools.
- 1 ½ Monte Carlo Productions completed (currently half way through second round).
- 1st Production: UK produced 25% of MC.
- O(60 billion events generated) so far in round 2
- Disk at GRID sites required only for staging (transfer data back to CNAF).

Of MICE and LHC

Sam and Jens

On behalf of MICE (in particular
Henry Nebrensky and David
Forrest)

MICE

- Muon Ionisation and Cooling Experiment
- Currently in the construction and simulation stages
- Use UK Grid Storage for simulation output storage and dissemination.
- Workflow incorporates LFC, lcg_utils toolset.
- lcg-cr, lfc-ls, lcg-cp
- Files are addressed by LFN

MICE requirements

- Prefer POSIX semantics in tools where possible.
- “it’s great when the commands mimic unix commands right down to options and flags”
- Wouldn’t mind improved performance, especially in lfc metadata operations (lfc-ls etc).
- Otherwise, currently happy.

MICE additional notes

- 30 TB total RAW
- Checksummed tarballs
- Using standard gLite tools
 - LFC, lcg-utils
 - Are they using FTS?
 - IC investigating PhEDEx
- Local (end user) download with browser
 - Need for HTTP support (at SURL level, not (just) TURL)
 - dCache and DPM both have HTTP support
 - StoRM and CASTOR don't
- Glasgow + IC + outside (Milan?)

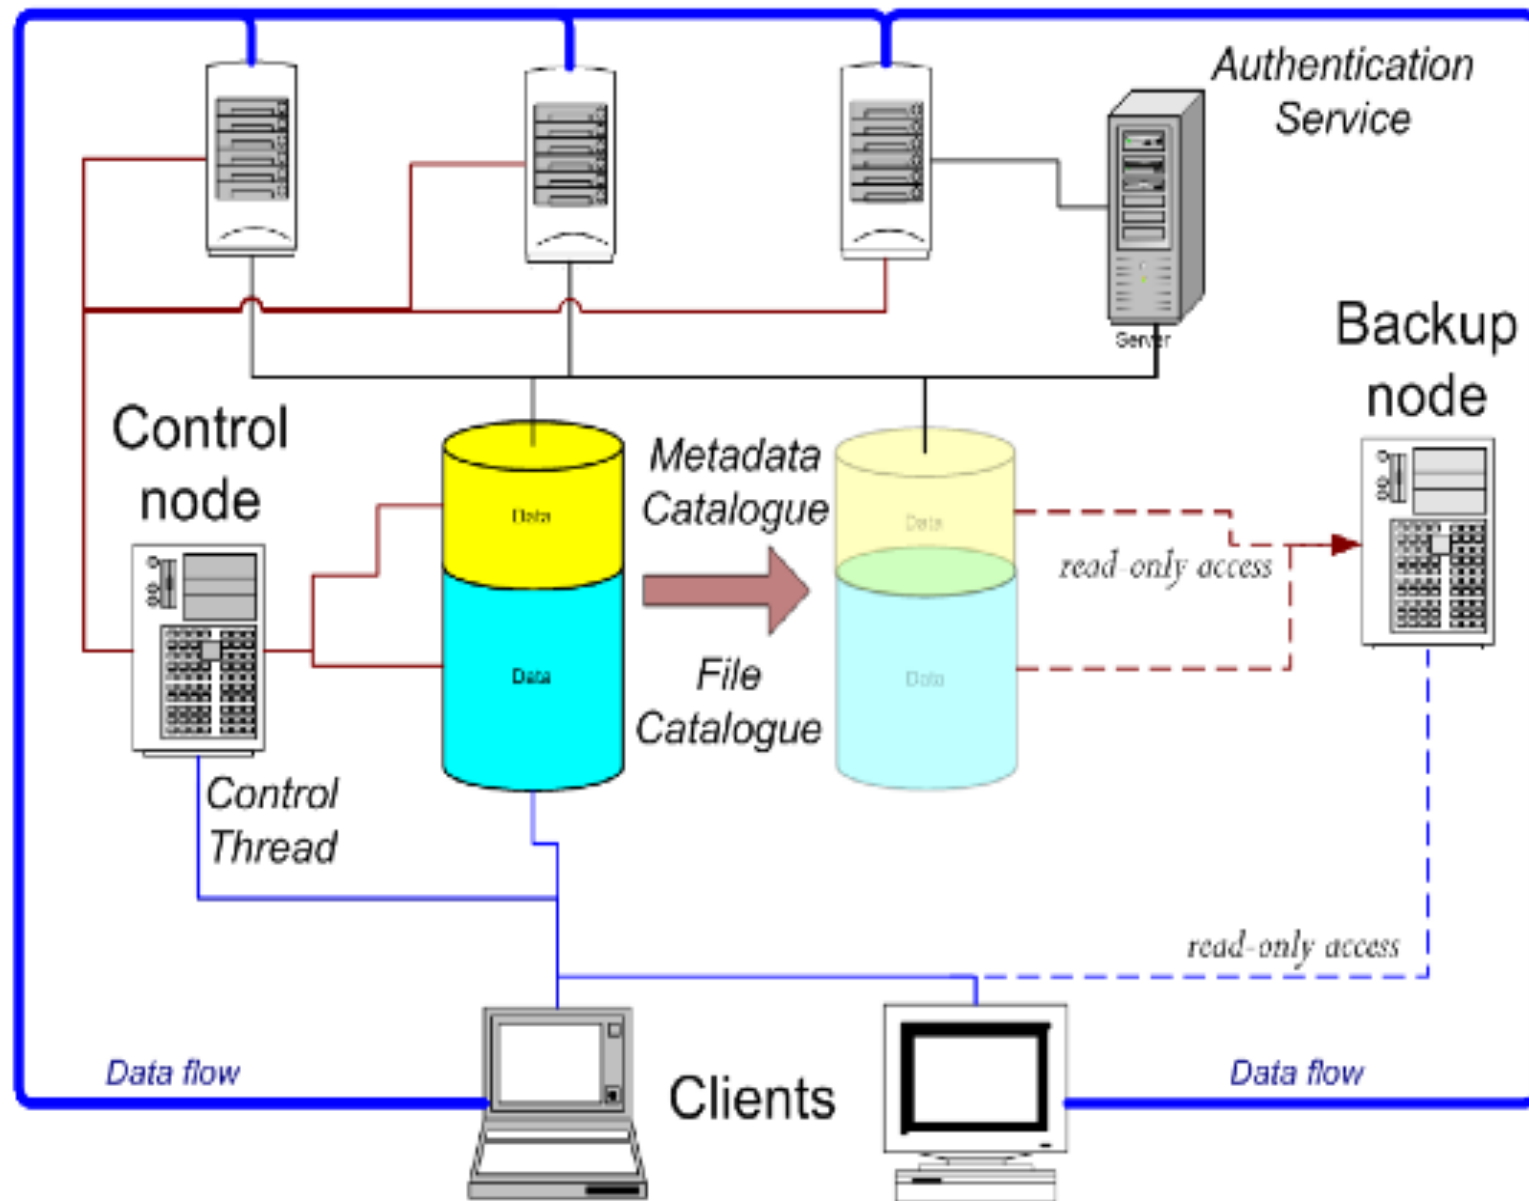
UKQCD/ILDG/DiGS and SRM/GridPP

Wahid Bhimji

Ukqcd / Digs

- QCD: Theory of the Strong nuclear force
- Computations done on space-time “lattice”
- Their data is collections of gauge configurations produced by super computers like HECToR
- These are stored on a data grid – “DiGS” designed by EPCC, Edinburgh - built on globus.
- Storage Elements use gridFTP but no SRM.

Storage Elements



Deployment architecture for DiGS data grid.

DiGS + SRM

- EPCCC recently added SRM client to DiGS.
 - Implemented own client with gSOAP and GSI.
 - (not lcg_utils as (e.g) they have own catalogue.)
- Sam and I helped them access GridPP resources at Glasgow and Edinburgh ECDF.
 - Tested with DPM and StoRM
- VO – ukqcd has a gridpp managed voms server but they are also a group in desy ildg voms
 - So using ildg voms for consistency

Camont

Matt, blindly distilling information
from Mark Slater and Karl Harrison.

Current Camont Storage Usage

- Camont have two main areas of Work - Web Spidering and Document Processing.
 - Web Spidering has no call for the use of Grid Storage and probably never will.
 - Document processing currently uses one of a pair of SEs for temporary (non-local) storage of job outputs before copying back to off grid disk.
 - This works for the current data sample of approximately 10^5 documents.
 - Each task is treated as a "perform-once" activity, so no incentive to pre-distribute documents.
 - But this probably wouldn't scale.

Camont Storage Future.

- To process larger numbers of documents, or increase the frequency of processing cycles "data management and job management would need to be decoupled".
- Sites would need to be pre-seeded with documents and output saved to the local SE (sound familiar?).
 - This poses an additional challenge as currently data is saved as tarballs which aren't readily usable by jobs.
- This would require of the order of a few hundred GB of stable storage at each site supporting Camont.
 - This could grow to TB-scale requirements after a few years.
- Camont is a perfect example of a VO that could benefit from the wisdom of the GridPP Storage Tribe.