

---

Grid2003 - the Project

Grid3 - a Functioning Demonstrator Grid -

Some Lessons Learned

Ruth Pordes

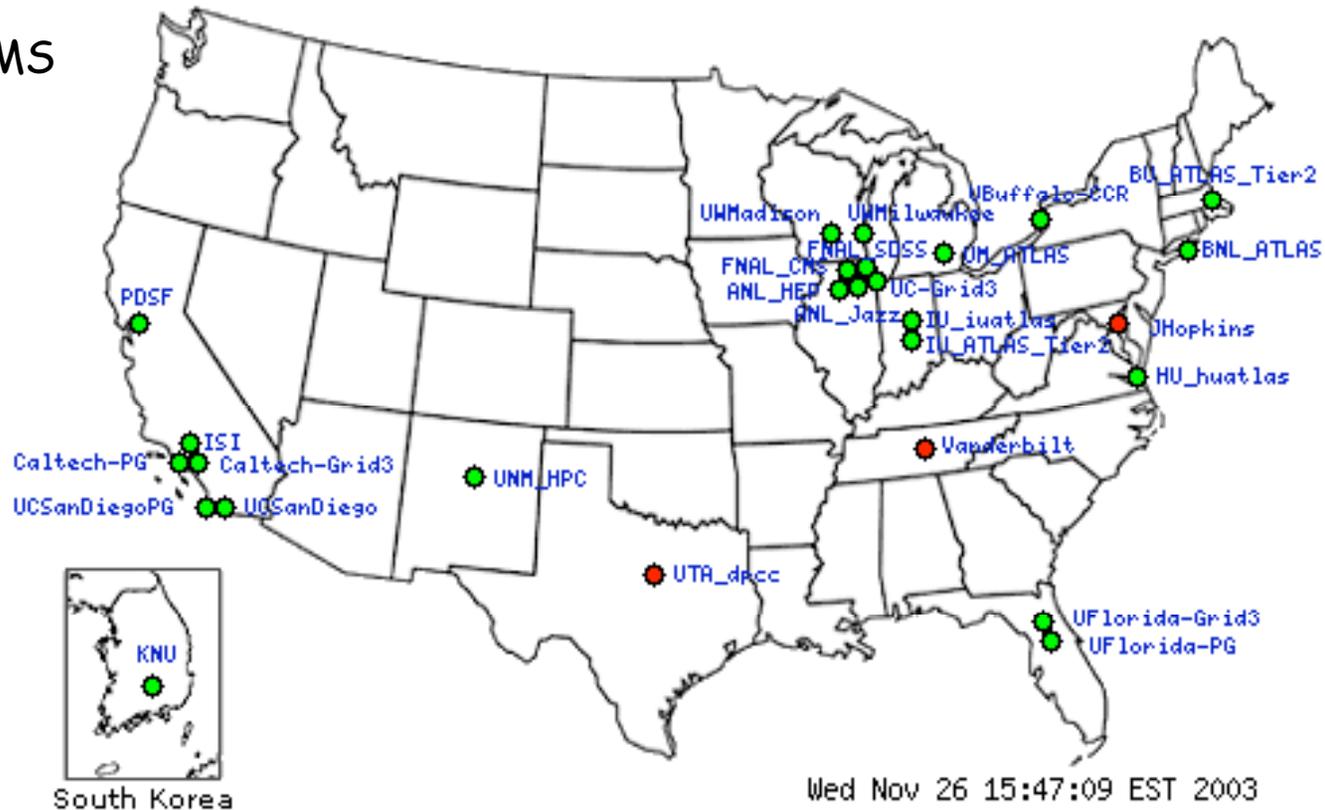
Fermilab

---



Grid3 is a continuation and extension of Existing Grid Testbeds..

- iVDGL-2 -> partnership with PPDG
- U.S. ATLAS
- U.S. CMS

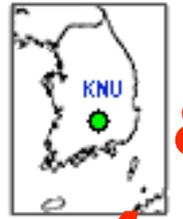


# Grid3 - A Federated Application Grid of Existing Resources



Grid3 is a continuation and extension of Existing Grid Testbeds..

- iVDGL-2 -> partnership with PPDG
- U.S. ATLAS
- U.S. CMS



South Korea

**8 Application Groups**  
**4 Virtual Organizations**  
**26 sites, 2500 CPUs**

Wed Nov 26 11:47:09 EST 2003

## Federated - Why that word for Grid3?



The LHC Single VO Testbeds have been successful with the "federated" approach

- local responsibility for facilities — but reporting to US LHC projects.
- systems and support, local resources w/ well defined interfaces.
- general grid-wide services provided by some of the sites.

The Sites - Autonomy, Control, Agreements, Policy, mix Local and Grid Use

- setup and manage the systems, install & configure the middleware Head node(s).
- interfaces well defined through VDT and services.
- automated central monitoring, validation, diagnosis to avoid scalability wall.

The Grid System and Services

- collaborative approach to bringing up cross-site services, e.g. VO management, monitoring, configuration management.
- interfaces well defined.
- robust against single point of failure.

The Applications

- Applications group responsible for end-to-end operations, diagnosis and the production services



## Grid2003 - the project



### Inclusive Coalition

- End to End Collaboration - Applications, Services, Sites
- U.S. ATLAS & U.S. CMS
- Trillium - iVDGL, PPDG, GriPhyN, brings in LIGO, SDSS, Computer Science Researchers
- Additional partners enabled to join during the project (Korea, Biology).
- 

### Project Lifecycle - beginning, middle, end

- Written and Endorsed Plan - July
- Execution and Milestones - July->November-> SC2003
- Lessons Learned and Next Steps, End Game Documents and Review - Dec, Jan
- 

### Team Approach

- Core + Taskforce + Stakeholders
- Principles + Plans + Decision Process + Post-mortem
- Inclusive Distributed Peer Collaboration
  - Flooded Mailists,
  - Weekly Team Meetings (Core, Taskforce, Steering)
  - Face to face workshops
  - Focussed Fests



## Inclusive Coalition

- End to End Collaboration - Applications, Services, Sites
- U.S. ATLAS & U.S. CMS
- Trillium - iVDGL, PPDG, GriPhyN, brings in LIGO, SDSS, Computer Science Researchers
- Additional partners enabled to join during the project (Korea, Biology).
- 

## Project Lifecycle - beginning, middle, end

- Written and Endorsed Plan - July
- Execution and Milestones - July->November-> SC2003
- Lessons Learned and Next Steps, End Game Documents and Review - Dec, Jan
- 

## Team Approach

- Core + Taskforce + Stakeholders
- Principles + Plans + Decision Process + Post-mortem
- Inclusive Distributed Peer Collaboration
  - Flooding Mailinglists,
  - Weekly Team Meetings (Core, Taskforce, Steering)
  - Face to face workshops
  - Focussed Fests

**Build and Operate a Common Grid Infrastructure, Share the Resources and Run Application to the Benefit of the Stakeholders.**



PPDG includes the running experiments and focusses on integrating Grid technologies step by step into the existing **Production Systems**.

Experience shows that the Grid Technologies are immature and evolving. A large scale Application Grid like Grid3 is not expected to be able to provide the **Quality of Service** required by experiments already in data taking mode.

Grid3 is a Demonstrator. Of course, since we have milestones we take meeting our metrics and success seriously.

Please keep this in mind when I get too enthusiastic about this project... the STAR Data Replicator, the Run II SAM-GRID are all production oriented efforts - they have to **Work for the End User Scientists**.



- Grid3 uses existing and emerging services. Given the aggressive milestones there was little time for development. Most developments was scripts and configuration automation. We should have done more!
- 
- Application Demonstrators and "Challenge Problems" were deployed on Grid3 after being demonstrated to run on the grid testbeds of the responsible groups.
- 
- Grid3 strategy is to interface with existing computing installations on a few well defined interfaces: batch queues, information service providers, and storage elements. to have as little impact as possible on the computing resources.
- 
- Bring the environment with whenever possible and assume as little as possible about the existing setup.
- 
- Applications must be able to interoperate and run on non-dedicated resources.
- 
- Applications must be able to install themselves dynamically, thereby imposing minimum requirements on grid facility managers.

## The Application Demonstrators



Blind search for continuous gravitational waves: LIGO analysis

DIAL: ATLAS analysis

GCE: USATLAS simulation

maxBcg: SDSS analysis

MOP: USCMS simulation

SnB: Bio-molecular analysis

BTeV simulation

GADU/Gnare: Genome analysis

Exerciser

GridFTP Demo

# Application & System Metrics



- related to LHC data challenge requirements for CMS and Atlas.
- metrics goal is to give quantities to aim for and a basis for discussion.

Metric (to be defined)	Target (to be defined)	Comments
Number of processors	up to 500	Possibly collected through MDS, archived and time-stamped
Data transferred per day	>2-3 TB	Data “flux” will need to be defined carefully.
Peak number of concurrent jobs	Up to 1000	Collect the total number of jobs running on Grid3, sorted by VO, archived so that time-dependent plots can be made.
Percentage of resources used	up to 90%	
Efficiency of job completion	up to 75%	Success to be defined. Other efficiency metrics could be identified.
Number of users	>10	Collect and sort by VO origin.
Number of different applications	>4	By an application “registry”, such as proposed by the WorldGrid “project” mechanism, this could be simplified.
Number of sites running multiple applications	>10	Collect with time stamps, intervals to be determined.
Rate of Faults/Crashes	<1/hour	Measure, and perhaps categorize by degree of severity.
Operational Support Load of full demonstrator	<2 FTEs	How many people (and in which roles) were operating Grid3?



## VDT - Middleware Packaging and Distribution Testing, Release and Acceptance Procedures and a Support Model



- Globus and Condor as packaged by the NSF Middleware Initiative (NMI);
- 
- iVDGL ATLAS developed software meta-packaging and distribution tool Pacman,
- 
- GriPhyN Chimera, Pegasus, DAGMAN Virtual Data System (VDS);
- 
- VO management scripts from the EDG;
- 
- Schema and information providers from the joint DataTAG/EDG/Trillium GLUE project;
- 
- The MonALISA monitoring framework from U.S. CMS, developed at Caltech;
- 
- The Netlogger monitoring package from the DOE Science Grid (DOESG) project.
-



Processing sites (Compute Elements) consist of a single head node with a gatekeeper service and a set of worker nodes.

All nodes have a shared file system, used for central access of applications, and input and output data transfer within a site.

- This is an incompatible configuration for the LCG-1 which does not allow a shared files system (They have made a wrapper for the job manager)
- Disk management needs work in next steps.

Supported batch systems through the GRAM/Globus job manager interface" LSF, PBS (including Maui extensions), FBSNG, Condor are all used.

No special software installations on the worker nodes: Ganglia broke this principle.

Multiple Application use of the Site Resources.

- Site Administrator maintains accounts; VOMS maintains mapping of User to Group
- Job execution policy w implemented through batch queue priorities based on VO group accounts.
- Policy that jobs of the VO owning the site have higher priority than those from other VOs . seems to work, but we have not done in depth analysis of the details.



Grid3 used the EDG/DataTAG Virtual Organization Management Software (VOMS) for central registration of user information. Probably first production use at scale.

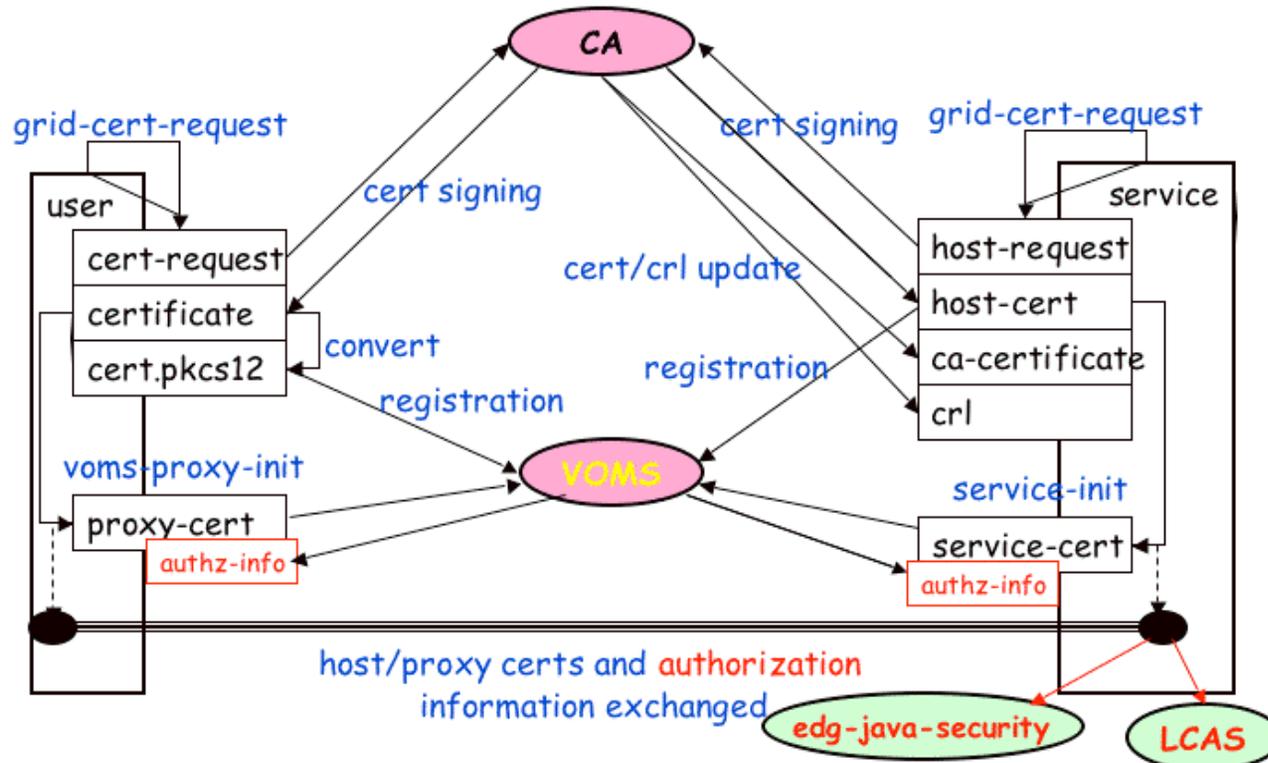
4 VOMS servers - 2 at Fermilab, 1 at BNL, 1 at Indiana.

Maps to group accounts at local site, where each account is associated with a specific VO.

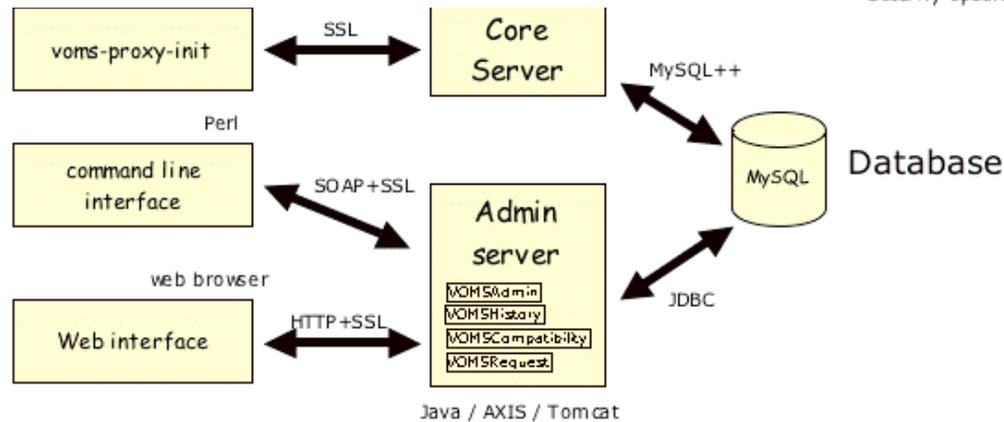
This short-term pragmatic approach allows job policy and monitoring based on a simple VO mapping.

Each site must run scripts calling administrative interface of VOMS server. Each host that does this must register its certificate with the VOMS database. This was not robust or scalable.

# User's Perspective 2.x



Security Update on 2.x - n° 4





Data movement services are provided through the use of GridFTP. Sustained transfers of four Terabyte of data per day across Grid2003 were demonstrated. Once initial firewall and open port issues were resolved, data movement has been stable across the Grid.

U.S. CMS has deployed a "Storage Element" implementation that provides a standard interface to the Grid based on the SRM interfaces. The site storage services at U.S. Tier-1 and Tier-2 centers are based on the dCACHE, while at CERN the LCG works on a Storage Element based on Castor, interfaced through the same SRM interface. This allows CMS to access and move data across the Tier-0 at CERN and the Tier-1 and Tier-2 sites as part of Grid2003.

U.S. ATLAS, LIGO and SDSS use the Globus/EDG Replica Location Service (RLS) in conjunction with job planning and workflow components of the GriPhyN Virtual Data



MDS is used as the information service. Sites publish information about the directories into which applications should be in-stalled, temporary data should be written, and input and output data should be staged. Information from the whole grid is available through a single top-level GIIS and associated web pages.

During the project we updated the MDS version, which required an update by all sites. Pacman supports this in an easy fashion, and the update was accomplished without perturbing running jobs. One lesson learned was that all such upgrades resulted in taking a lot of effort for configuration and tuning, highlighting need for more planning, automation and testing for future configuration management and tuning.

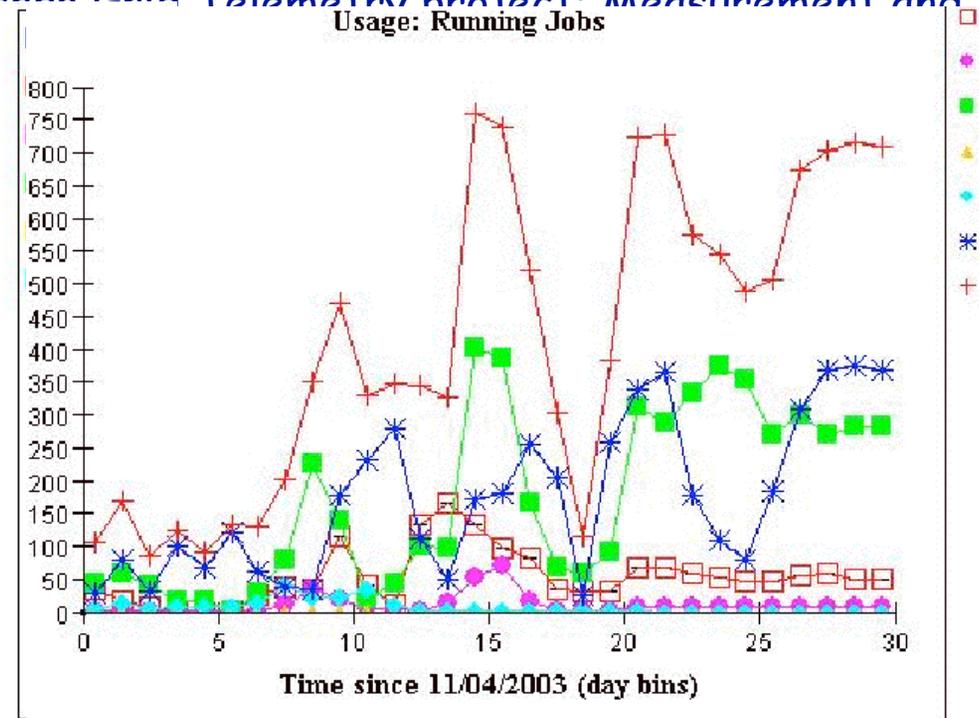
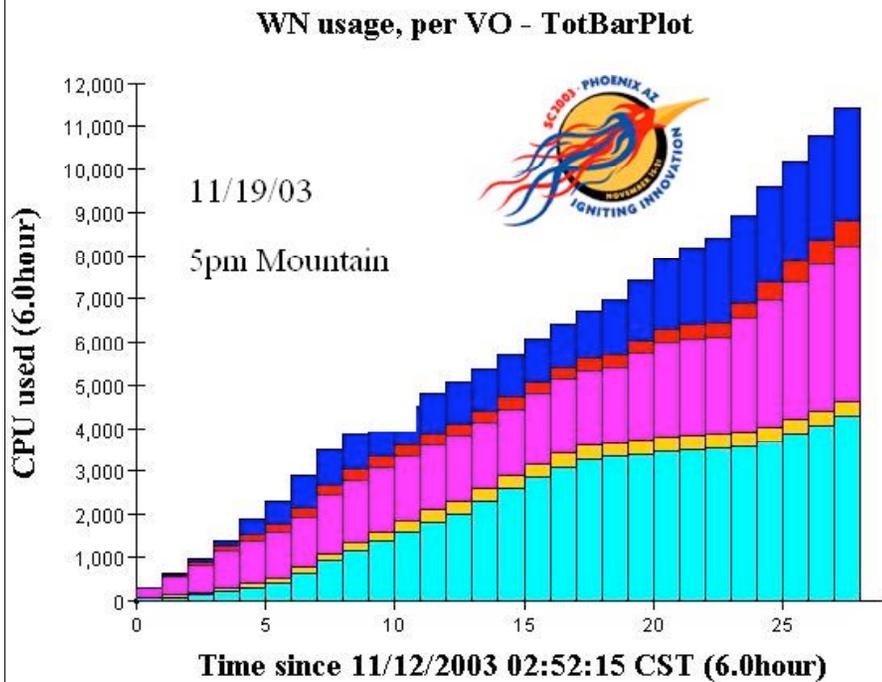
In general we spent inordinate amounts of time on configuration issues and fixing the same or similar problems. Any time we upgraded or changed a core service it took much effort and many emails and phone calls to get back to a stably running system.



The MonALISA system was used as the core monitoring system. It was extended for Grid2003 to collect and present information based on each VO. This allowed accounting for usage of resources by each VO on each Grid2003 site.

MDS, MonaLisa, Ganglia, MDViewer and ACDC (the University of Buffalo monitoring database and job manager monitoring system) were all variously used for monitoring. Each of these monitoring applications has a different scope. All were useful.

Collaboration with the University of Chicago Grid Telemetry project: Measurement and





The iVDGL Grid Operations Center (iGOC) based at Indiana University contributed central Operations Services and provided support for central servers: GIIS, MonaLisa and VOMS databases.

In general the distributed model for operations, contacts and support worked for this project.

Identifying coordinators for various services helped in organization and enabled delegation of responsibilities to help in the 'scaling' problem as the size of the project grew, but clearly more work is necessary here to understand operations support at larger scales.

The site administrators were responsive, and responsive to phone calls if needed when the actual deadlines were approaching.

The use of email lists ensured maximum communication, but the volume of mail was overwhelming.

# The Catalog



<b>● IU_ATLAS_Tier2</b>	<a href="http://atlas.iu.edu">atlas.iu.edu</a>	CS	IN	USATLAS	Allen, Matt	64
Authentication:	GRAM Authentication test failure: authentication failed: GSS Major Status: Authentication Failed GSS Minor Status Error Chain: init.c:497: globus_gss_assist_init_sec_context_async: Error during context initialization init_sec_context.c:204: gss_init_sec_context: SSLv3 handshake problems globus_i_gsi_gss_utils.c:848: globus_i_gsi_gss_handshake: SSLv3 handshake problems: Couldn't do ssl handshake OpenSSL Error: s3_pkt.c:1046: in library: SSL routines, function SSL3_READ_BYTES: sslv3 alert certificate expired				2003-12-8 00:19:33 GMT	
Hello World:	UNKNOWN					2003-12-8 00:19:34 GMT
Long Job:						
-Batch Query:	UNKNOWN					2003-12-8 00:00:00 GMT
-Batch Sub:	UNKNOWN					2003-12-8 00:19:34 GMT
-Batch Cancel:	UNKNOWN					2003-12-8 00:00:00 GMT
gsiftp:	UNKNOWN					2003-12-8 00:00:00 GMT
<b>● IU_iuatlas</b>	<a href="http://iuatlas.physics.indiana.edu">iuatlas.physics.indiana.edu</a>	CS	IN	USATLAS	Luehring, Fred	4
Authentication:	GRAM Authentication test failure: connecting to the job manager failed. Possible reasons: job terminated, invalid job contact, network problems, ...				2003-12-8 00:00:02 GMT	
Hello World:	UNKNOWN					2003-12-8 00:00:02 GMT
Long Job:						
-Batch Query:	UNKNOWN					2003-12-8 00:00:00 GMT
-Batch Sub:	UNKNOWN					2003-12-8 00:00:02 GMT
-Batch Cancel:	UNKNOWN					2003-12-8 00:00:00 GMT
gsiftp:	UNKNOWN					2003-12-8 00:00:00 GMT
<b>● KNU</b>	<a href="http://cluster28.knu.ac.kr">cluster28.knu.ac.kr</a>	CS	Korea	USCMS	Cho, Kihyeon	3
Authentication:	Pass					2003-12-8 00:06:00 GMT
Hello World:	Pass					2003-12-8 00:06:33 GMT
Long Job:						
-Batch Query:	Pass					2003-12-8 00:06:07 GMT
-Batch Sub:	Pass					2003-12-8 00:06:02 GMT
-Batch Cancel:	Pass					2003-12-8 00:06:08 GMT
gsiftp:	Pass					2003-12-8 00:06:44 GMT



Interoperability tests with LCG going on now and at RSIS conference:

- Storage Management and Data Movement
- Jobs
- 

Many of the working areas in Grid3 are already joint projects between the LCG and Trillium or the S&C projects.

- Glue Schema
- Testing
- Packaging
- Additional collaboration in areas of monitoring and operations have been discussed over the past few months.

And the PNPA-RG ... what is that .



Interoperability tests with LCG going on now and at RSIS conference:

- Storage Management and Data Movement
- Jobs
- 

Many of the working areas in Grid3 are already joint projects between the LCG and Trillium or the S&C projects.

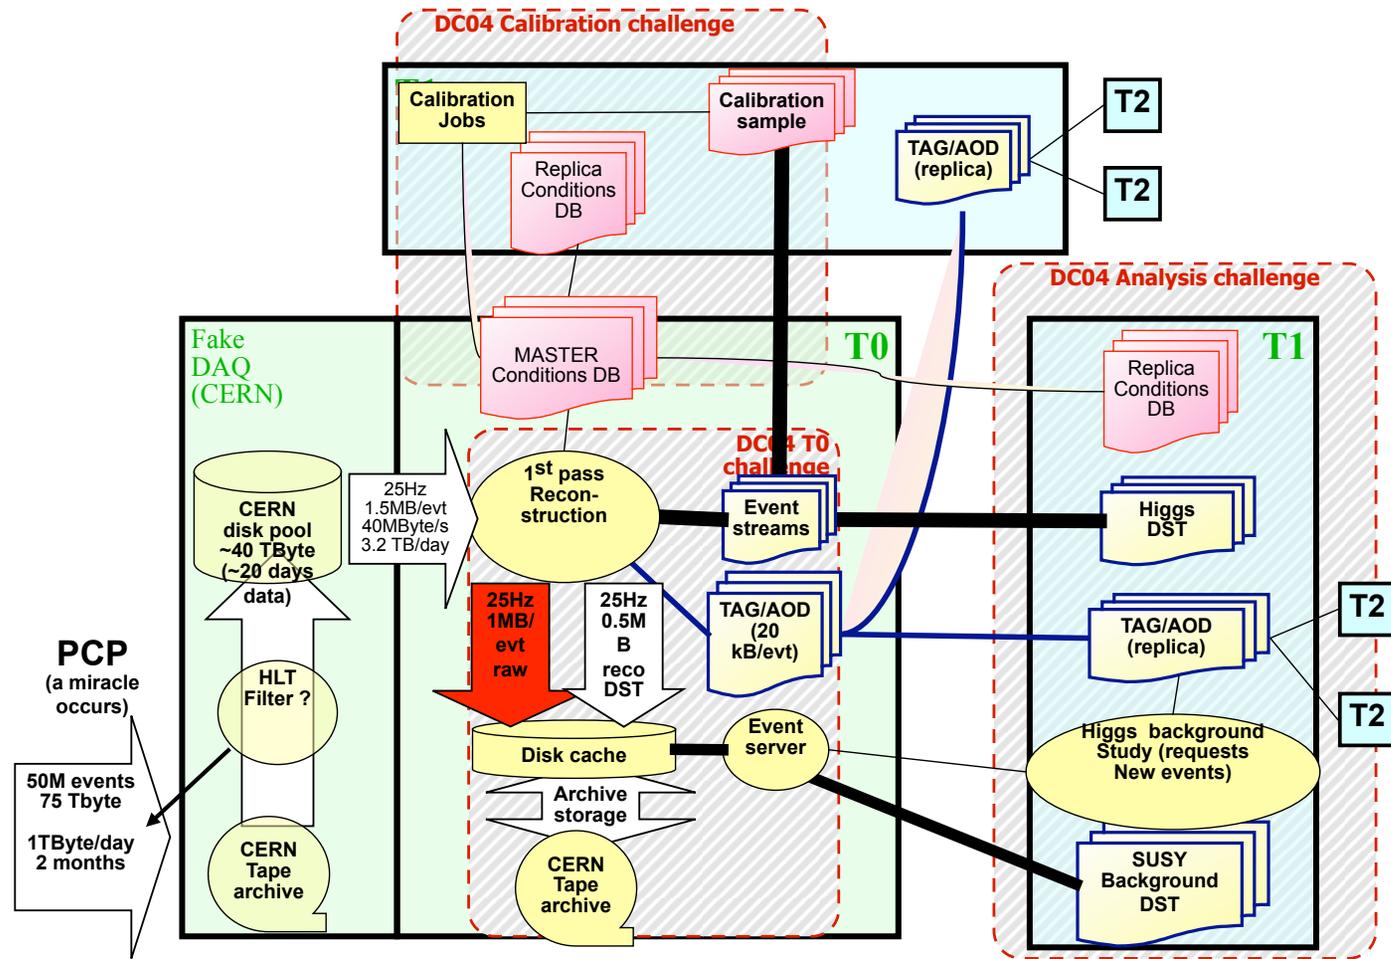
- Glue Schema
- Testing
- Packaging
- Additional collaboration in areas of monitoring and operations have been discussed over the past few months.

And the PNF A LCG ... what is that?

**We Will Interoperate with the LCG  
today, tomorrow and and the next day!**



Grid2003 will be used for US DC04 analysis challenge.  
 Production will continue to use Grid3





- 1) Monte Carlo production
  - if possible extending to non-U.S. sites, using Chimera-based tools
- 2) Collect and archive MC data at BNL
  - MAGDA, RLS/RLI/RLRCs involved
- 3) Push MC data files to CERN
  - MAGDA
- 4) Reconstruction at CERN
  - using LHC/EDG components
- 5) MAGDA "spider" finds new reconstruction output in RLI and copies them to BNL
  - may require interfacing to EDG RLS
- 6) Data reduction at BNL, creating "collections" or "datasets" and skimming out n-tuples
  - DIAL, maybe Chimera
- 7) Analysis: Distributed analysis of collections and datasets.
  - DIAL, GANGA



PPDG's focus on End to End Application Grids has proven its worth.

The next phase of PPDG plan is to work on a common grid infrastructure between the DOE Laboratories and focus on Facility-Grid interfaces and Production Services.

Propose a Common Project involving all Labs with effort from the Experiment Teams, coordinated by an expanded Executive Team.

Strengthen the Partnerships with iVDGL, DOE SG etc.

For the Shared Application Grid - ie Grid3++ :

- Not ready for prime time.
- Expect to see additional services e.g. Storage and Data Management
- New Experiments and Applications - Run II? RHIC? Life Science? NCSA PDQ
- Continued focus on tools and automation
- Step towards more Proceduralized Operational Support
- Continuous and increased interoperability with the LCG

...Grid4, OSG-0, Cats&Dogs will interoperate (federate) with LCG-2.x, -3.y...



Grid3 is a first step which gives us a platform to move to become  
"Open Science Grid-0"



A National Grid Infrastructure for Science:

Open Science Grid (<http://www.opensciencegrid.org>)

Grid3 demonstrates that this not just a pipe dream - we can turn it into reality if we try.

Vicky White, Head of Fermilab CD - "Grid2003 is not a fancy grid, it's only a first step, but it does what a grid is supposed to do for scientific research. It successfully combines computing power from multiple sources and transports data to and from offsite locations to allow individual scientists in different experiments to share computers that are not under their control and solve scientific problems.



Initial commitment is to provided the U.S. LHC grid resources as a scalable, engineered and managed PRODUCTION grid SERVICE.

This is a long way from Grid3 which is running in the Hero support mode:



Tim Thomas notices at xx that his authentication does not work.

The GRid3 catalog page shows ~10 sites failed at this time and that the error message is the same as Tim got (~2 sites have different errors messages).

- **Those who make the Catalog Page are Heroes.** They made a page that a) shows the actual errors encountered; b) shows at a glance the scale of the problem; c) shows at a glance the problem crosses SITES and VOs. but that many sites are not affected. d) is very usable in being able to select options and see information from all sites on one page.

Ed May finds that his VDT 1.1.12 site has no CRLs and works and his VDT 1.1.11 site has CRLs and does not work.

- **Those who ensured that there be test sites and a test grid are Heroes,** and that there are the same site administrator for both the production and test sites so we have many knowledge individuals with the ability to run local tests and checks

A Google search on the error message pretty much confirms that the problem is due to an updated CRL. This meshes with many sites who do not do automatic updates not seeing the problem.



Tim comes in on Sunday, learns more and identifies the ESNET ROOT CRL expired within minutes of him noticing a failure:

Certificate Revocation List (CRL):

Version 1 (0x0)

Signature Algorithm: sha1WithRSAEncryption

Issuer: /DC=net/DC=ES/O=ESnet/OU=Certificate Authorities/CN=ESnet Root CA 1

Last Update: Dec 6 02:13:49 2002 GMT

Next Update: Dec 6 02:13:49 2003 GMT

**Tim is The Hero. He came in at the beginning and followed it through to the end.**



Grid2003 is a Distributed Collaboration involving an End to End Team representing and empowering Applications, Services, and Sites.

There is distributed ownership, distributed responsibility and distributed commitment.

The fact that Grid3 is Multi-VO , that we have tools that everyone shares, that we have mail lists that reach everyone, helped in diagnosing and identifying a problem that would have affected the whole community (come Monday for most).

We know we can't run production grid operations this way for the LHC and certainly neither RHIC nor RUn II can rely on a service - but such projects show us how to do it and will ensure our success .



Grid2003 is a Distributed Collaboration involving an End to End Team representing and empowering Applications, Services, and Sites.

There is distributed ownership, distributed responsibility and distributed commitment.

The fact that Grid3 is Multi-VO , that we have tools that everyone shares, that we have mail lists that reach everyone, helped in diagnosing and identifying a problem that would have affected the whole community (come Monday for most).

We know we can't run production grid operations this way for the LHC and certainly neither RHIC nor RUn II can rely on a service - but such projects show us how to do it and will ensure our success

**Thank you to All -**  
**It was a good shift**