


Grid2003 Activities

Ian Fisk

October 16, 2003




Grid
Projects

Experiments



What is Grid2003?



Grid2003 is a project to create a multi-VO Grid infrastructure in the US


- ➔ Collaborative effort of Trillium Grid Projects, US-ATLAS, US-CMS, LIGO, SDSS, and BTeV
- ➔ Natural extension of current testbed activities
- ➔ Pilot implementation of a persistent, national grid infrastructure.

Grid2003 is primarily an integration and deployment project

- ➔ The schedule and manpower do not permit a lot of development



Grid2003 Goals




Grid2003 is a project, done by a collaboration with well-defined stakeholders

- ➔ contributions from the US ATLAS and US CMS S&C projects, LIGO, SDSS, BTeV and US-Grid Projects: iVDGL, GriphyN and PPDG

The purpose of the Grid2003 is a project to build a grid environment to:

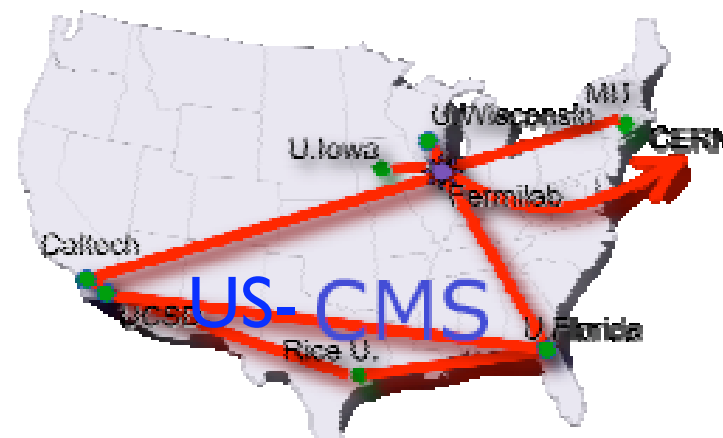
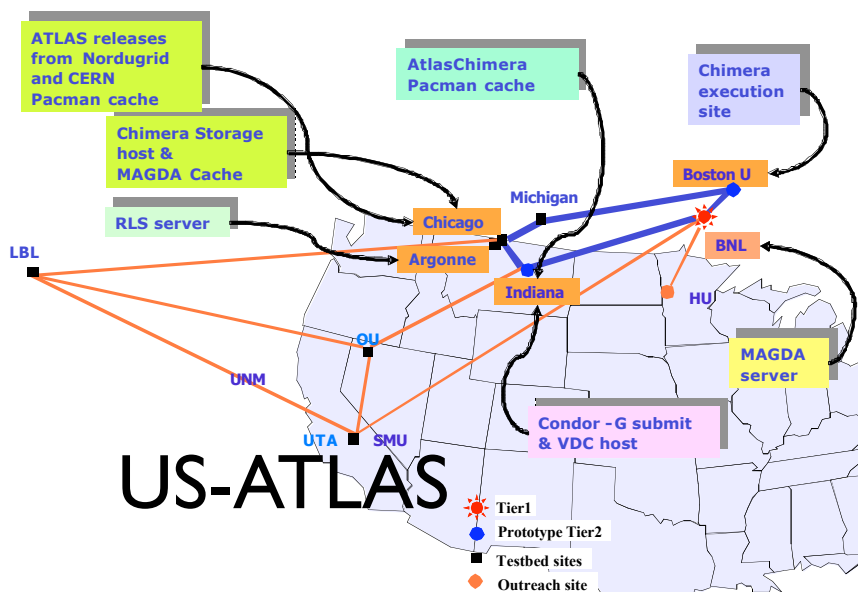
- ➔ Provide the next phase of the iVDGL Laboratory
- ➔ Provide the infrastructure and services needed for LHC production and analysis applications running at scale in a common grid environment
- ➔ Provide a platform for computer science technology demonstrators
Provide a common grid environment for LIGO and SDSS applications, and possibly others




US-ATLAS and US-CMS both have well developed testbed activities

- Rolling prototypes for development
- Facilities at reasonable scales
- Successfully used for official simulated event production

Up to now these have been single VO installations





Worked toward achieving a well defined set of metrics for success

- ➔ Definitions of functionality and scale

Define the minimum set of services needed on sites to support the VOs

- ➔ Simple services and common installations
 - Rely heavily on VDT (Virtual Data Toolkit)
- ➔ Make as few requirements as possible on the sites themselves
 - Only put services on gateway systems
 - Separate privilege requiring installations from non-privileged
 - Should be no requirements except OS version on Worker Nodes

No components go into Grid2003 that haven't been evaluated and tested on individual VO testbeds


- ➔ Multi-VO environment and new scale is hard enough, without having to debug component functionality



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 450	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?



Grid2003 Work Flow



Grid2003 Flash Workflow





Currently all services are installed on a gateway node


- ➔ Privileged access is required for primary middleware and parts of experiment middleware
- ➔ Currently the Grid2003 site installation is 5 steps
 1. Set up local scheduler
 - Many sites use Condor, but we have installations of PBS, LSF, and FBSNG
 2. Get Host certificate
 3. Create Local Accounts
 - We are currently using generic VO accounts (uscms01,usatlas1,ligo1, sdss1, btev)
 4. Get pacman version
 5. Install Grid2003 packages with pacman
 - Download, installation, and configuration is currently handled by pacman
- ➔ To install a new site should be very fast, but we are currently averaging a few hours of work on each site to get the configuration right



- ➔ Central Grid2003 cache is based on VDT 1.1.11
- ➔ All applications are using dynamic installations
- ➔ All Grid2003 sites are publishing to VO GIIIS, which in turn report to Grid2003 GIIIS
 - Successfully moved to MDS 2.4
 - Grid2003 extension of the GLUE schema to provide applications with the location of application, data, and temporary directories on the participating sites
- ➔ VOMS servers for the VOs are on-line
 - Ran into complexity issues but persevered
 - The fact that all sites need to exchange hostcerts with all VOMS servers is complicated as the number of sites and VOs grows
- ➔ MonaLisa and Ganglia both reporting
 - Configuration not as automated as the rest of the software



Current Status of Grid2003 Facilities



There are more than 25 Grid2003 sites.


Currently there are approximately 2500 CPUs reporting to Grid2003

- ➔ About two thirds of these are shared resources
 - Grid2003 users don't get all of them, but we've hit full utilization

All sites are publishing to the GUIS



Current Status of Grid2003 Usage



We have over 50 people registered as Grid2003 users

- ➔ Though a smaller number are actively submitting jobs

We seen over 1000 simultaneous grid submitted jobs

We also seen nearly 100% utilization of the grid resources

- ➔ A combination of grid submitted jobs and locally submitted jobs


Combination of data analysis jobs, detector simulation jobs, computer science demonstrators all running together



We've transferred over 4TB of data in a single day using GridFTP transfers between sites

We implemented a SRM (Storage Resource Manager) based Storage Element at a few sites

- ➔ These provide a uniform interface to storage systems
- ➔ Along with advanced functionality
 - In addition to point-to-point transfers, there is the capability of
 - resource reservation
 - transfer queuing and prioritization
 - third party initiated transfers

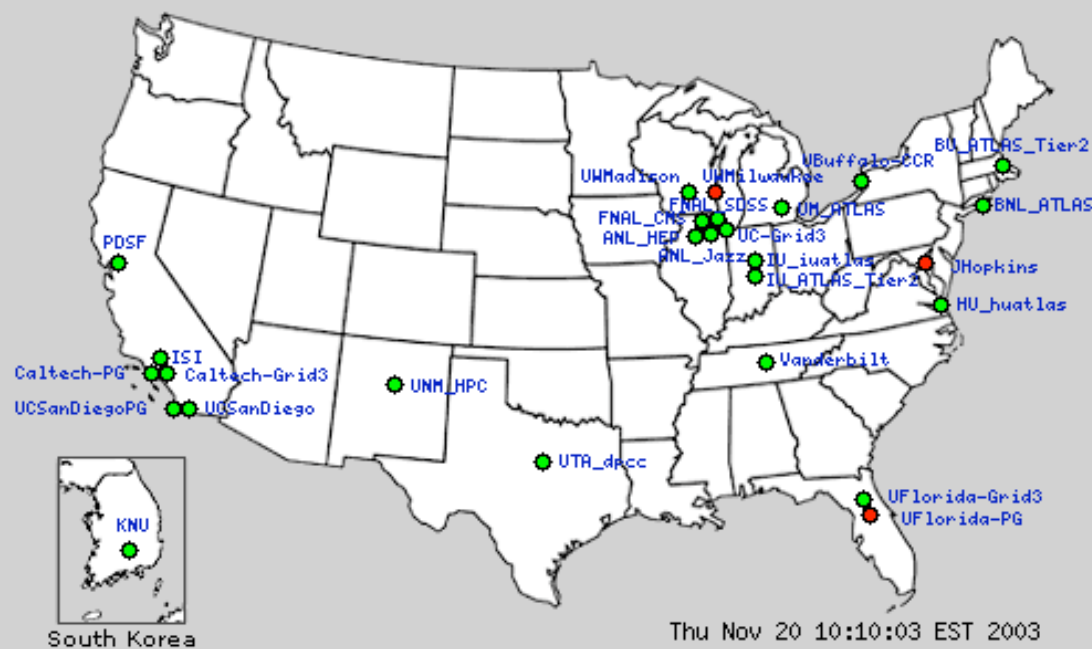


Grid2003 has more than met it's metric for success

We see Grid2003 as a pilot toward a US national grid infrastructure called the Open Science Grid

Current Site Status

Click on a site for detailed information





We are proposing to DOE and NSF a roadmap for the U.S. to build a national grid infrastructure for science, the **Open Science Grid (OSG)**

We propose a program of work to federate the U.S. LHC grid resources into a scalable, engineered and managed grid, the Open Science Grid.

- ➔ Much of the LHC grid infrastructure in Europe will be provided by a combination of CERN central resources and a consortium of European centers that propose to federate some of their resources in a grid for e-science in Europe EGEE. Centers in the U.S. and other parts of the world will federate with this European grid infrastructure in order to provide the global computing grid for LHC science.
- ➔ It is now time for the U.S. to also federate its LHC computing resources and in doing so to continue to lead the efforts towards a global grid infrastructure for LHC science and beyond. We propose to provide and operate these resources at the national laboratories and universities as the initial seed for the Open Science Grid

This new Peta-scale computational service will be built as an open national infrastructure, optimizing shared use of resources for diverse collaborative research. The Open Science Grid will serve as a backbone to merge grid computing efforts of allied experiments in particle and nuclear physics, and can be extended to other scientific communities