Strategic Plan for Grids (2006-2008)

Eileen Berman, Amber Boehnlein, Ian Fisk, Ashutosh Kotwal, Ruth Pordes
6 page version. 9/30/06.
The Grid Strategic plan is designed to cover two years. These years represent a unique challenge for the Computing Division (CD) to support the Run II experiments with increasing luminosity and contracting effort, assist CMS during the critical first years at the new energy frontier, facilitate the migration to core grid services of the existing programs of FNAL, and to maintain flexibility to actively participate in ILC preparations.

Mission

To support (develop, innovate, operate) and inform the scientific community of FNAL in the use of Grid technologies.

To partner with others in developing and deploying Grid middleware, taking a leadership role to ensure effective production grid infrastructures, including the Open Science Grid.

To provide common services at FNAL, including the operation of the FermiGrid campus grid.

Context and Assessment of Current State

The experimental community is making good use of Grid Resources. CDF, CMS and DØ generate large MC samples, and DØ reconstructed 0.8B events on the Grid. Experiments such as MiniBooNE are using grid resources via FermiGrid routinely.

Integrated production grid facilities depend on Facilities, Networking, Storage and Data Movement, Lab, and Scientific Core Services to provide the underlying middleware and infrastructure. FNAL offers a production grid facility (FermiGrid) which provides a set of common services and a portal to Fermilab compute and storage services. Stakeholder interoperability and resource sharing are a goal. Near term plans involve increasing the scalability, availability and throughput. Offsite network bandwidth will soon support six 10GE production network links. This is believed to be sufficient for the next few years. Storage resources are driven by the needs of the scientific community of FNAL. Planned capacity for data storage and movement will meet these needs for the next few years. A longer-term strategy for more general use by additional grid users needs development.

A production grid is also dependent on production middleware. CD is actively involved in grid middleware development, integration and deployment for authorization and identity management, accounting, resource and workload management, data handling, and storage and data movement. Awareness and application of good security practices and policy is increasing. Strengthening of the grid infrastructure to production quality is necessary also. Regular meetings provide feedback on plans, status, and immediate issues. Fermilab personnel spend dedicated effort triaging problems and insuring that the offered services do

work. This effort is ongoing but would benefit greatly from a longer-term strategy where effort is dedicated to this task. Many people are involved in OSG planning, deployment, and operation.

Vision

By 2008 FNAL will be seamlessly integrated into the distributed computing system for CMS as a Tier-1 center and as an analysis facility. Processing resources for event reconstruction, data selection, and regional analysis will be transparently accessible through the Open Science Grid infrastructure. Custodial data storage from the experiment and event simulation centers, as well as data serving to remote analysis facilities will be accessible through reliable and secure common grid interfaces.

By 2008 the Run II experiments will have completed the migration to stable operations with a set of core computing services that perform at the scale required to process and analyze the final years of data taking. Computing services will be migrated to common grid solutions where this is seen as the most efficient program of work, weighing the development cost against the potential savings in operations and support. In areas where legacy solutions are operated the transition to maintenance will include the identification of sufficient effort to provide adequate support for the Run II program.

By 2008 FNAL will be the preeminent computing facility in the Open Science Grid (OSG). FNAL will provide an advanced set of core services that provide resource management, auditing, accounting and virtual organization management services at FNAL. OSG communities will have access to and discovery of opportunistic computing resources at FNAL. The work with the extension projects and the storage implementation in the Virtual Data Toolkit (VDT) will secure FNAL, in collaboration with DESY, as the primary provider for grid enabled storage solutions for the OSG. FNAL will have sufficient effort and expertise to facilitate and guide the transition of smaller communities to common grid services and components where desired.

Stakeholders

The sponsors of the grid work are the Fermilab Computing Division base program and the DOE and NSF in the several year SciDAC-2 and NSF sponsored projects. Effort and deliverables are provided from many groups both internal and external to the Laboratory.
The customers are:
The running experiments based at FNAL (CDF, D0, Minos, MiniBooNE) and the astrophysics community.
CMS.
Simulation and theory including accelerator modeling, ILC, LQCD.
The Computing Division in its pursuit of excellent, effective facilities.
The members of the Open Science Grid Consortium.

Goals and Objectives

The Computing Division will provide development, deployment, and operations of common grid services and components that will enable many communities to work efficiently in distributed computing environments. These efforts are identifiable in the contributions of FermiGrid, the many common projects often in the context of Open Science Grid, and the direct computing division contributions to the major stakeholders.

FermiGrid will continue to provide core services that are needed by local communities to facilitate adoption and deployment of grid enabled computing. Additionally FermiGrid will develop and deploy common gateway services to enable more of the FNAL computing resources to contribute to the OSG.

FNAL will contribute to the success of the Open Science Grid in the areas of development of grid services, in the contribution of facilities and resources, and equally importantly in the aiding of communities to adopt common grid solutions. FNAL will actively participate in the area of authorization, auditing, resource management and accounting, integration and scale testing, as well as grid-enabled storage services through the Virtual Data Toolkit (VDT). FNAL will work with smaller communities to try to ensure their needs are satisfied by the developing grid infrastructure.

As the current major computing project of CD, CMS has a unique position in driving the grid computing activities. FNAL is dedicated to assisting CMS in achieving the functional, scale and reliability of grid services needed to complete the distributed computing infrastructure as outlined in the CMS Computing Technical Design Report (CTDR).

FNAL is committed to the completion of the Run II program. The objectives of the Run II program are to reach a stable and sustainable level of operations in an era when the integrated data set doubles yearly while the collaborations contract in size. The concentration will be streamlining operations and implementing solutions that can be operated in the environment of diminishing effort. CDF and DØ will continue to generate Monte Carlo data using grid resources. The DØ collaboration will continue to use grid resources to support collider data reconstruction. Both collaborations expect to use FermiGrid to support analysis computing. CDF plans to migrate all its analysis computing to FermiGrid. The Computing Division is also committed to the success of the neutrino and astrophysics programs, working with them to leverage grid resources to achieve their computing needs.

Computing Division goals include mitigating the risk of disruption for the Run II experiments as grid services are replaced or augmented. The Run II collaborations are expected to provide and support a high level user interface and experiment specific layer of services. The design goal is that these layers should rely on common interfaces and grow thinner over time as grid services mature. The Run II experiments understand that underlying common services can and will be replaced in order to meet CMS goals and objectives, and that there is limited or no support for services outside of common areas or

non-standard use or deployment of common services.

As the requirements of the ILC community evolve over the next two years, FNAL will maintain sufficient expertise and flexibility to actively contribute to the program. The current accelerator and detector modeling efforts need reasonably modest resources, but this will grow over time.

Strategies

1. CD will provide a gateway to a common pool of CPU resources, through FermiGrid. This strategy consolidates the system support effort and also maximizes the efficiency of CPU and storage utilization.

It is desirable for all experiments to access FermiGrid via a core set of robust Grid middleware tools and services with well-defined functionalities. These will then have a broad user base. The basic functionality provided by these tools should be maintained by the Grid community - developers and users at Fermilab and elsewhere. These tools and services should be called or accessed by the experiment-specific software through a relatively thin interface, such that the interface is transparent and as easy to maintain as possible. The software should also be designed and used in such a way that as the underlying implementation evolves, the experiments' interfaces either do not have to change at all, or the changes required are minimal and easy to implement and validate.

It is desirable for the tools and services used to access FermiGrid to also provide the basic functionality needed by the experiments for offsite resources. The experimental interfaces for FermiGrid vs offsite resources should be as similar as possible.

2. FNAL contributes to the OSG consortium and project for the operation, support, and evolution of a common, shared, distributed facility across the US (and in other countries where there are members of the consortium). Responsibilities include project management and planning, security and authorization and storage services (integration and deployment of SRM/dCache), technical writing and communication. Contributions to the OSG software stack help ensure that the common core grid services are usable by and appropriate for the FNAL experiments and scientific communities. CD makes its computing and storage resources accessible to the OSG and enables opportunistic use of the resources by other OSG communities. FNAL contributes to the development of new services with“External Projects” that the OSG relies on.

OSG contributions include work in interoperation between OSG services and equivalent services on EGEE and TeraGrid, in support of the global data distribution, storage, processing and analysis systems of the FNAL scientific communities.

3. . The US CMS software and computing project includes a grid services and interfaces area that works with the distributed computing tools and tier-1 facility areas to deliver the

US CMS distributed facility. CMS will be operating its globally distributed system in production from the end of CSA06, (November 2006) throughout the ramp up of the system to support commisioning and data taking. CMS has defined a set of " baseline grid services" which meet its needs for data storage, distribution, processing and analysis of event and simulation data. It relies on this core grid middleware being supplied and supported by external grid technology groups, the CD and other identified projects, through the lifetime of the experiment.

CMS relies on and contributes to the FNAL production grid facility and the OSG distributed facility. The collaboration interfaces and integrates its higher level data management, distribution and processing services on top of multiple grid infrastructures – in particular the EGEE and OSG. US CMS institutions with computing and storage resources are members of the OSG. CMS encourages common services across EGEE and OSG to the extent possible. CMS contributes to the development of new services with"External Projects" that the OSG relies on, as part of Joint Projects at FNAL with Run II and CD, and participates in the Worldwide LHC Computing Grid project to ensure the development of new services in EGEE and OSG are appropriate to its needs.

4.To support the Run II experiments CD will continue to provide support for the existing critical computing infrastructure:

DØ relies on SAMGrid job management for job submission for all production activities (including FNAL DØ farm production) and is actively transitioning to a forwarding mechanism to use LCG and OSG mechanisms for brokering, resource selection and job execution. Remote collaborators are hosting SAM stations for data delivery to grid farms. SAMGrid enters maintenance in June 2007 after completion of development projects to adapt to standard grid interfaces and services and provide additional monitoring capability.

CD will support SAM data handling operations for CDF and DØ with a combined team, and will support needed underlying storage services interfaced to SAM via the Storage Resource Manager interface. CD will support SAMGrid job submission.

CDF is currently using a grid facility called NAmCAF, which relies on Condor glide-ins and other Condor tools to support job submission and execution. NAmCAF provides a single point of job submission for CDF users, using the existing user interface, and directs the jobs to available CPU resources at OSG sites (currently five and increasing) across North America. CDF expects to use NAmCAF as its primary resource for Monte Carlo production in North America. This facility is supported by the CDF collaboration.

Run II strategies include having well-trained support teams and defining and monitoring metrics in order to anticipate any operational or scaling problems and to identify and

mitigate any problems which do occur as quickly as possible. An additional strategy is a yearly operational plan written by the collaboration management, with estimated resource needs and timetables for unusual activities.

In the cases when the objectives of scalable systems, reliance on standard components used through common interfaces, and reducing effort for stable operations have not yet been achieved, the collaboration and CD will outline a specific strategic plan which includes incremental tactical steps towards the final objectives within the resources available.

Resource Needs

After June of 2008 all resources purchased for the LHC will likely be entirely consumed with LHC activities, so there should not be any planned reliance on securing large numbers of opportunistic computing resources after this date.

CMS effort will contribute to both the experiment specific services and to the common grid services and infrastructures, increasingly shifting to operations from development over the next 2 years. The pool of developers dedicated to Run II support will be transitioning over the next year into development of common services.

As Grid services and production is integrated into the Scientific and Core Facilities at the lab and experiments, explicit Grid developments will transition to service development for resource management.

Progress Indicators

Development projects milestones will be tracked and reported.
For experiment layers, compliance with standards will be tracked.
The following will be reported monthly, taking account of the experiment needs:

-- The % of jobs through grid interfaces per experiment.
-- The % of data through grid interfaces per experiment.
-- Number of sites offering turnkey grid access per experiment.
-- Efficiency of use of grid resources for Monte Carlo production.
-- Numbers of problems a week and the successful resolution of said problems.

Additional Information

The risks of this strategic plan include:
The continued support and attention of the underlying Grid technology groups – Condor and Globus – to deliver the standard middleware used by all grid infrastructures.
Succces in the annual review for funding of the Open Science Grid.
Failures in communication with the experiments to ensure their ongoing requirements and plans are met.