BUILDING THE U.S. NATIONAL FUSION GRID: RESULTS FROM THE NATIONAL FUSION COLLABORATORY PROJECT

D.P. Schissel,[1] J.R. Burruss, [1] A. Finkelstein,[2] S.M. Flanagan, [1] I.T. Foster,[3] T.W. Fredian,[4] M.J. Greenwald,[4] C.R. Johnson,[5] K. Keahey,[3] S.A. Klasky,[2] K. Li,[2] D.C. McCune,[2] M. Papka,[3] Q. Peng,[1] L. Randerson,[2] A. Sanderson,[5] J. Stillerman,[4] R. Stevens,[3] M.R. Thompson,[6] and G. Wallace [2]

[1]General Atomics, P.O. Box 85608, San Diego, California 92186-5608
email: schissel@fusion.gat.com, Phone: (858) 455-3387, Fax: (858) 4156
[2]Princeton Plasma Physics Laboratory, Princeton, New Jersey 08543
[3]Argonne National Laboratory, Argonne, Illinois 60439-4844
[4]Massachusetts Institute of Technology, Cambridge, Massachusetts 02139-4307
[5]University of Utah, Salt Lake City, Utah 84112
[6]Lawrence Berkeley National Laboratory, Berkeley, California 04720

The U.S. National Fusion Collaboratory Project (http://www.fusiongrid.org/) is developing a persistent infrastructure to enable scientific collaboration for all aspects of magnetic fusion research. The project is creating a robust, user–friendly collaborative software environment and making it available to the more than one thousand fusion scientists in forty institutions who perform magnetic fusion research in the United States. In particular, the project is developing and deploying a national Fusion Energy Sciences Grid (FusionGrid) that is a system for secure sharing of computation, visualization, and data resources over the Internet. The goal of FusionGrid is to allow scientists at remote sites to participate fully in experimental and computational activities as if they were working at a common site thereby creating a virtual organization of the U.S. Fusion community. The project is funded by the USDOE Office of Science, Scientific Discovery through Advanced Computing (SciDAC) Initiative and unites fusion and computer science researchers to directly address these challenges. This paper reports on the services presently deployed on FusionGrid and discusses the future needs of the fusion research community.

1. Introduction

Grid computing refers to the large–scale integration of computer systems via high speed networks and advanced middleware to provide on–demand access to data analysis capabilities and related functions not available to one individual or group of machines. Using shared languages and interaction protocols, grid systems reach out across the globe to access the computing resources, information and services required to satisfy local user needs. To users, the highly integrated networks that embody grid systems are transparent so that services furnished from afar appear to be provided by local computers. This concept is illustrated in present day life by the electrical power grid. When a user plugs a device into an electrical wall socket, the how and where of electrical generation are immaterial. Useful functions (electricity) are hidden by an interface (plug) that conceals the details (power plant) of how they are implemented allowing the individual to solely concentrate on using the function. Thus, computer Grid technology has the capability to enable large–scale scientific and business collaboration among members of virtual organizations, remote experimentation, and high–performance distributed computing and data analysis.

The National Fusion Collaboratory Project [1] is funded by the United States Department of Energy (DOE) under the Scientific Discovery through Advanced Computing Initiative (SciDAC) to deploy a pilot Grid or Collaboratory called FusionGrid for the U.S. magnetic fusion energy research community. The vision for FusionGrid is that experimental and simulation data, computer codes, analysis routines, visualization tools, and remote collaboration tools are to be thought of as network services (in the analogy above electricity is a network service) which represents a fundamental paradigm shift for the fusion community. This Grid will create a virtual organization of the U.S. fusion community whose resources are protected by a shared security infrastructure including strong authentication to identify users and fine–grain authorization to allow stakeholders to control their own resources. FusionGrid will shield the users from software implementation details and allow a sharper focus on the physics with transparency and ease–of–use being the crucial elements. In this environment, access to services is stressed rather than data or software portability. FusionGrid is not focused on computer cycle scavenging (e.g. SETI@home) or distributed supercomputing that are typical justifications for Grid computing.

The implementation of FusionGrid is a collaboration itself since the project unites fusion scientists and computer scientists from seven institutions to form a closely coordinated team. This group is leveraging existing computer science technology where possible and extending and/or creating new capabilities where required.

Goals. The goal of FusionGrid is to allow scientists at geographically distributed sites to participate fully in experimental and computational activities as if they were working at a common site. The ultimate expression of this goal would be that a scientist working in a windowless office would not be able to tell their geographic location since all of the services they need for their research would be seamlessly available via FusionGrid.

Accomplishing this goal will advance scientific understanding and innovation in magnetic fusion research by enabling more efficient use of existing experimental facilities and more effective integration of experiment, theory, and modeling. Physics productivity will be increased by (1) enabling more efficient utilization of experimental time on the three large facilities through more powerful between pulse data analysis resulting in a greater number of experiments at less cost, (2) allowing more transparent and uniform access to analysis and simulation codes, to data, and to visualization tools resulting in more researchers having access to more resources, (3) creating a standard tool set for remote data access, security, and visualization allowing more researchers to build these into their own services, (4) facilitating the comparison of theory and experiment, and (5) facilitating multi–institution collaborations. The Project will also increase the productivity of code and tool developers by (1) supporting more users with fewer installations at reduced cost, (2) facilitating shared code development projects resulting in more rapid code creation, and (3) creating a standard tool set for remote data access, security, and visualization allowing these services to be easily built into new code systems.

FusionGrid Authentication Security. The Internet is an open system, where the identity of the communicating partners is not easy to ensure. Furthermore, the communication path traverses an indeterminate set of routing hosts and may include any number of eavesdropping and active interference possibilities. Thus, Internet communication is much like anonymous postcards, which are answered by anonymous recipients. For Internet based commerce the solution is to use encryption (to assure privacy and security) and certification (to assure that communication is happening between desired endpoints). Security on FusionGrid focuses on the latter, thus assuring that any communication occurs between known endpoints.

FusionGrid security employs Public Key Infrastructure (PKI) to secure communication on the Internet through the use of a public and private cryptographic key pair that is obtained and shared through a trusted authority. When using a key pair, only one of the keys, referred to as the private key, must be kept secret. The other key, referred to as the public key, can be disseminated freely for use by any person who wishes to participate in a secure transaction with the person holding the private key. This is possible because the keys are mathematically related but it remains computationally infeasible to derive the private key from knowledge of the public key.

FusionGrid uses the X.509 certificate standard and the DOE Grids CA (Certificate Authority) to implement PKI for secure communication. The CA uses an X.509 certificate to bind a public key to the name contained in the certificate and assures third parties that this binding is valid for both name and key. A scientist who desires to join FusionGrid will generate a public/private key pair and apply to the DOE Grids CA for an X.509 certificate binding that public key with their name. The CA will verify their identity and the validity of their request (i.e. determine they are known members of the community and they have a reason to join FusionGrid). As required by FusionGrid, the private key is password protected on the user's computer.

The implementation of authentication on the FusionGrid is accomplished using the Globus Toolkit [2]. To log into FusionGrid, All a user needs to do is issue the *grid–proxy–init* command once per day and type in the password for their private key. This single sign–on is accomplished behind the scenes by the use of a short–lived proxy certificate that is derived from the user's long–term X.509 certificate. The proxy certificate uses its own unencrypted private key, so that is can make frequent authenticated connections on behalf of the user to multiple services without requiring additional password interactions with the user. The benefit to the user is that they need only log–on once no matter how many different services they desire to use.

FusionGrid Authorization Security. Once a user's identity has been validated they are still not given open access to all FusionGrid resources (codes, computers, visualization tools or data). These are made available only to those users who have the proper authorization to use them. The authorizations are granted and controlled by the stakeholders who provides the resources and not by any central FusionGrid authority. The combination of strong authentication and authorization is seen in every day life with the commercial airplane industry in the U.S. When using a commercial airplane, a government ID must be obtained once prior to the journey to verify the passenger's identity when boarding the plane. But verifying identify is not enough to use the resource. Authorization for usage of the resource is required in the form of a boarding pass and the verification of that pass as they step onto the plane.

In FusionGrid, the Akenti authorization service [3] is used to control access to services. Akenti is an established authorization service designed to make access decisions for distributed resources controlled by multiple stakeholders. When an authorization decision needs to be made, the Akenti policy engine gathers all the relevant information for the authenticated user and resource, and determines the users' rights with respect to the resource.

FusionGrid Data Access.. Data access on FusionGrid has been made available using the MDSPlus data acquisition and data management system [4] combined with the relational database Microsoft SQL server. MDSplus is by far the most widely used and successful data acquisition and management system in the international fusion community. Presently it is used at more than 30 sites spread over 4 continents and has become the de facto standard for data access.

MDSplus and the Globus Toolkit have been combined to create secure X.509 certificate based client/server data access on FusionGrid using the standard MDSplus interface without any loss in speed or functionality. Presently, the three main MDSplus experimental data repositories at Alcator C–Mod, DIII–D, and NSTX are securely available on FusionGrid. Data management by MDSplus of large datasets generated by simulation codes is presently being tested using results from NIMROD simulations. NIMROD is a 3D MHD simulation code that runs on very large parallel computers [4]. Using the MDSplus server at DIII–D, output from NIMROD runs up to 100 GB have been stored and served to users for both further data analysis and scientific visualization. Secure access to SQL Server has been done via MDSplus since the integration of Globus and Microsoft SQL Server is not yet complete.

FusionGrid Computational Codes. The code TRANSP, used for time dependent analysis and simulation of tokamak plasmas [4], was released as a service on FusionGrid late in 2002. Running on a Linux cluster at PPPL, over 900 TRANSP runs from seven different experimental machines have been completed within the FusionGrid infrastructure. This approach has drastically reduced the efforts to support and maintain the code which were previously required of the developers and by sites using the code. TRANSP analysis by Murakami using FusionGrid was presented at the 2002 APS/DPP meeting [5].

Details on the implementation of codes as services on FusionGrid, and TRANSP specifically, are covered in a companion paper by Burruss. The general sequence for a scientist to use a code service on FusionGrid is to first prepare and store a code’s input data in an MDSplus repository. When ready, a code start command with a pointer to the input data is issued and once the run has completed the output data is written to the same MDSplus repository. During code execution a FusionGrid monitoring system (FGM: see Flanagan) is provided so that the user can follow the entire computation process.

For ease of usage, a GUI utility called PreTRANSP was created for input data preparation and TRANSP code invocation. Prior to using PreTRANSP the user needs to log into FusionGrid only once even though they are using at least three different services (MDSplus, FGM, TRANSP).

Adding new codes as services to FusionGrid is relatively straightforward. The scientist providing the service needs to make decisions related to input and output data as well as where the service will run with the basic FusionGrid infrastructure handling all other requirements.

FusionGrid Visualization Services. The SCIRun Problem Solving Environment (PSE) [4] is being used on FusionGrid for an advanced scientific visualization service. SCIRun is open source software that gives the user the ability to create complex 3D visualizations. SCIRun on FusionGrid gives the scientist a no–license–cost visualization capability that rivals present day commercial visualization packages. Hardware requirements for SCIRun are modest, basically a mid–range PC running Windows or Linux with a mid–range graphics card is sufficient. Initial usage of SCIRun has been with NIMROD simulation data with visualization results including animations being presented by Brennan in an invited talk at the 2002 APS/DPP meeting [5].

8. FusionGrid Advanced Collaboration Services. The goals of FusionGrid's advanced collaborative environment service is to use computer mediated communications techniques to enhance work environments, to enable increased productivity for collaborative work, and to exploit the use of high–performance computing technologies to improve the effectiveness of large–scale collaborative work environments. Examples of such collaboration include off–site support of experimental operations, large group collaborations in a tokamak control room, simulation/experimental data analysis meetings, and shared code debugging.

Tiled display walls are being used to enhance the collaborative work environment of the tokamak control room [4]. As its name implies, a tiled display wall utilizes multiple projectors tiled together to build a bright, high–resolution, seamless display. Utilizing commodity based hardware and computer cluster technology, tiled walls of arbitrary size and resolution can be constructed at modest cost. Today displays that are 16’ x 8’ with 20 million pixels are not uncommon. Such a display offers a large–format environment for presenting either a high–resolution visualization or more multi–source smaller visualizations to a collaborative group than would be possible on standard computer display screens.

As a prototype FusionGrid service, tiled display walls have been tested in a variety of usage modalities ranging from two tiled walls geographically separated being tied together by software to form shared collaborative displays to a single tiled wall used for collocated group sharing and discussion. As currently envisioned, the initial tiled wall service for FusionGrid will be the installation of a 2–tile system in the NSTX and DIII–D control rooms. The service will allow any researcher either in the control room or off–site, with proper authentication and authorization, to share any X–windows based visualization, with the entire control room. For scientists within the control room, this interactive shared visualization takes the place of “passing around” a graphical printout or “calling over” scientists to collaboratively view a normal desktop display. For scientists off–site, this service gives them the capability to interactively share visualizations and participate in experiments, something previously not possible.

The Access Grid [4] is used by FusionGrid to create a service that enables group–to–group interaction and collaboration that improves the user experience significantly beyond teleconferencing. The Access Grid includes the ability to utilize for scientific research a complex multi–site visual and collaborative experience integrated with high–end visualization environments including tiled display walls. Developed exclusively for a FusionGrid service, the personal interface to the Access Grid (PIG) has been developed as a low cost alternative to a full–blown conference room size Access Grid node. Presently PIGs have been installed at C–Mod and DIII–D and a full–scale AG node at PPPL. Usage so far has been for meetings and working group sessions that would normally have been done by teleconference. In the near future, a joint experiment between C–Mod and DIII–D will use PIGs in the control rooms for joint experimental analysis.

Future Directions. The main thrust of future work for FusionGrid is aimed towards ease–of–use including easier certificate management, easier and more uniform access into FusonGrid code services possibly via web portals, and better interplay between FusionGrid security and site security (firewalls). At the same time, new code services will be deployed including the microturbulence code GS2 running on a Linux cluster at the University of Maryland and the stability code Gyro running at the Topical Computing Facility at PPPL. For FusionGrid's advanced collaboration service the challenge is to develop a science of collaboration; to understand what works and why (human factors and social issues), what functionality is needed, and how should it be interfaced to the users. The upcoming tests of both tiled walls and Access Grid nodes will help to solidify the requirements for FusionGrid's advanced collaboration environment through real world testing.

Acknowledgment

Work supported by US Department of Energy under Cooperative Agreement No. DE-FG02-90ER54084.

References

[1] http://www.fusiongrid.org/.

[2] http://www.globus.org/.

[3] http://www-itg.lbl.gov/Akenti/homepage.html.

[4] http://www.fusiongrid.org/services/.

[5] http://www.aps.org/meet/DPP02/.

FIGURE CAPTION

Fig. 1. FusionGrid will provide a unified framework in which data, analysis codes, visualization tools, and remote collaboration services, are available securely and transparently over the Internet.

Fig. 2. (a) Tiled display wall being prototyped for a tokamak control room. (b) Personal interface to the Access Grid used for remote collaboration.