

GA-A25540
Rev. 1

A NATIONAL COLLABORATORY TO ADVANCE THE SCIENCE OF HIGH TEMPERATURE PLASMA PHYSICS FOR MAGNETIC FUSION

**FINAL REPORT TO THE
U.S. DEPARTMENT OF ENERGY
for the period
AUGUST 1, 2001 THROUGH JULY 31, 2006**

by

**D.P. SCHISSEL, G. ABLA, J.R. BURRUSS, E. FEIBUSH, T.W. FREDIAN,
M.M. GOODE, M.J. GREENWALD, K. KEAHEY, T. LEGGETT, K. LI,
D.C. McCUNE, M.E. PAPKA, L. RANDERSON, A. SANDERSON,
J. STILLERMAN, M.R. THOMPSON, T. URAM, G. WALLACE**

DATE PUBLISHED: DECEMBER 2012



DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

GA-A25540
Rev. 1

A NATIONAL COLLABORATORY TO ADVANCE THE SCIENCE OF HIGH TEMPERATURE PLASMA PHYSICS FOR MAGNETIC FUSION

FINAL REPORT TO THE
U.S. DEPARTMENT OF ENERGY
for the period
AUGUST 1, 2001 THROUGH JULY 31, 2006

by

D.P. SCHISSEL, G. ABLA, J.R. BURRUSS, E. FEIBUSH,* T.W. FREDIAN,[†]
M.M. GOODE,[‡] M.J. GREENWALD,[†] K. KEAHEY,[‡] T. LEGGETT,[‡] K. LI,[¶]
D.C. McCUNE,* M.E. PAPKA,[‡] L. RANDERSON,* A. SANDERSON,[§]
J. STILLERMAN,[†] M.R. THOMPSON,[‡] T. URAM,[‡] G. WALLACE[¶]

*Princeton Plasma Physics Laboratory, Princeton, New Jersey

[†]Massachusetts Institute of Technology, Cambridge, Massachusetts

[‡]Lawrence Berkeley National Laboratory, Berkeley, California

[†]Argonne National Laboratory, Argonne, Illinois

[¶]Princeton University, Princeton, New Jersey

[§]University of Utah, Salt Lake City, Utah

Work supported by the
U.S. Department of Energy under
DE-FC02-01ER25455, DE-AC02-76CH03073,
DE-FG02-04ER54235, W-31-109-ENG-38,
DE-FC02-01ER25454 and DE-AC03-76SF00098

GENERAL ATOMICS PROJECT 30106
DATE PUBLISHED: DECEMBER 2012



1. INTRODUCTION AND BACKGROUND

This report summarizes the work of the National Fusion Collaboratory (NFC) Project [1,2] funded by the United States Department of Energy (DOE) under the Scientific Discovery through Advanced Computing Program (SciDAC) [3] to develop a persistent infrastructure to enable scientific collaboration for magnetic fusion research. A five year project that was initiated in 2001, it built on the past collaborative work performed within the U.S. fusion community and added the component of computer science research done with the USDOE Office of Science, Office of Advanced Scientific Computer Research. The project was a collaboration itself uniting fusion scientists from General Atomics, MIT, and PPPL and computer scientists from ANL, LBNL, Princeton University, and the University of Utah to form a coordinated team. The group leveraged existing computer science technology where possible and extended or created new capabilities where required.

Developing a reliable energy system that is economically and environmentally sustainable is the long-term goal of Fusion Energy Science (FES) research [4]. In the U.S., FES experimental research is centered at three large facilities with a replacement value of over \$1B. As these experiments have increased in size and complexity, there has been a concurrent growth in the number and importance of collaborations among large groups at the experimental sites and smaller groups located nationwide. Teaming with the experimental community is a theoretical and simulation community whose efforts range from applied analysis of experimental data to fundamental theory (e.g., realistic nonlinear 3D plasma models) that run on massively parallel computers. Looking toward the future, the large-scale experiments needed for FES research are staffed by correspondingly large, globally dispersed teams. The fusion program will be increasingly oriented toward the International Thermonuclear Experimental Reactor (ITER) [5] where even now, a decade before operation begins, a large portion of national program efforts are organized around coordinated efforts to develop promising operational scenarios. Substantial efforts to develop integrated plasma modeling codes are also underway in the U.S., Europe and Japan. As a result of the highly collaborative nature of FES research, the community is facing new and unique challenges. While FES has a significant track record for developing and exploiting remote collaborations, with such large investments at stake, there is a clear need to improve the integration and reach of available tools.

The NFC Project was initiated to address these challenges by creating and deploying collaborative software tools. The original objective of the NFC project was to develop and deploy a national FES “Grid” (FusionGrid) that would be a system for secure sharing of computation, visualization, and data resources over the Internet. The goal of FusionGrid

was to allow scientists at remote sites to participate as fully in experiments and computational activities as if they were working on site thereby creating a unified virtual organization of the geographically dispersed U.S. fusion community. The vision for FusionGrid was that experimental and simulation data, computer codes, analysis routines, visualization tools, and remote collaboration tools are to be thought of as network services. In this model, an application service provider (ASP) provides and maintains software resources as well as the necessary hardware resources. The project would create a robust, user-friendly collaborative software environment and make it available to the US FES community. This Grid's resources would be protected by a shared security infrastructure including strong authentication to identify users and authorization to allow stakeholders to control their own resources. In this environment, access to services is stressed rather than data or software portability.

2. PROJECT SUMMARY

Deployment of the production FusionGrid infrastructure (Fig. 1) has been made possible through work supported by the NFC project, by leveraging other SciDAC programs, and with the base funding for FES research [6]. Access to fusion data has been made secure via the integration of MDSplus [7,8] and the Globus toolkit [9,10]. The TRANSP code [11] has been released as a FusionGrid computational service along with supporting infrastructure development (data storage, monitoring, user GUI). This FusionGrid service has been so successful that it has become **the** production system for TRANSP usage in the United States, is widely used internationally [12], and was the model for numerous other fusion codes to be deployed on FusionGrid. The development of the desktop-based Access Grid (AG) node combined with the concept of application sharing in the AG venue [13,14] has allowed usage of this capability in the tokamak control room. Shared display walls [15] combined with new application sharing software has been deployed in tokamak control rooms with such positive results that fusion programs decided to purchase and permanently install new display and computer hardware to fully deploy this new functionality. SCIRun visualization software [16] has been integrated with MDSplus storage of large fusion simulation datasets to facilitate comparison of experiment and simulation. Finally, interacting with FusionGrid’s security system was greatly simplified by the deployment of MyProxy to manage scientists’ credentials as well as through a new web-interface-based centralized authorization system.

2.1. FUSIONGRID SECURITY

One of the major goals of the FusionGrid was to provide transparent and secure access to all of its resources. To implement this goal a unified FusionGrid authentication and authorization system was established based on X.509 identity certificates for authentication, the Grid Security Infrastructure (GSI) to transmit the authentication transparently and a shared authorization system. Originally, the X.509 identity certificates were issued by the DOE Grids Certificate Authority (CA) and the Globus toolkit (GT2) was employed to provide secure authentication, remote job execution, and communication over the open Internet. In response to the difficulties that fusion scientists experienced in managing their own X.509 credentials, the FusionGrid established a centralized Credential Manager that facilitates the issuing, use, renewal, and revocation of certificates [17]. The Credential Manager presents a customized Web interface to users to request membership in the FusionGrid. The user needs to provide name and contact information and select a FusionGrid username and password. The Credential Manager requests an X.509 credential from the FusionGrid CA. If it is approved and issued, it is stored on a central server and the




Fig. 1. Timeline of milestones from the fusion scientists' perspective accomplished by the SciDAC funded National Fusion Collaboratory Project.

user is added to the ROAM authorization user base (see below). The user can then use the ROAM web interfaces to request access to FusionGrid services. The Credential Manager supports MyProxy interfaces to retrieve proxy credentials for Grid service invocations, and PubCookie interfaces to allow single-signon to a number of FusionGrid web sites. The user only needs to login once with a username and password to establish 12-hour secure FusionGrid session. The centralized server provides administrators with a single place to view all the FusionGrid users and to easily revoke any user's privileges. The changes have been well received by users and administrators alike.

Since the operating policy of the DOEGrids CA did not allow third parties to possess user's private keys, the Credential Manger scheme required a different CA. ESnet agreed to run a custom CA for the FusionGrid that allows centralized storage of credentials. The Credential Manager is run at LBNL and mirrored at MIT to provide the required robustness in case of system or network failures. The scripts and servers that get or renew proxy certificates will try to contact the LBNL server first and if that fails contact the MIT server. The Credential Manager also provides a mechanism for renewal of proxy certificates. When a user gets a proxy certificate, he can specify that it is to be used for a long running job on a certain server. Then that server can renew the proxy, if it is going to expire before the job is completed. This feature is currently being used for TRANSP jobs.

The Credential Manger also supports PubCookie [18] interfaces to provide single sign-on to a set of FusionGrid web sites. These sites can use the PubCookies just for user authentication or to get a proxy certificate for the user so they can act as a portal for grid services. So far this feature has only been used for a limited number of web sites and a single test service. A prototype version of the Credential Manager that supports the use one-time-passwords (CryptoCard tokens) for login was also implemented. There is ongoing work to use the production Credential Manager to provide SecurID access to FusionGrid.

In response to the difficulty of installing a complete GT2 environment, required to use any FusionGrid resource, the NFC project repackaged the Globus libraries used for FusionGrid logins and secure MDSplus communications. The core FusionGrid software utilizes a very small portion of the complete Globus software package. By identifying which modules were actually being referenced by the FusionGrid software a “GsiLite” library was constructed using around 100 Globus source modules. A simple makefile was constructed to build this single library and the FusionGrid code was re-linked with this library. Initial work was successfully completed on the Linux platform. The GsiLite library was successfully built on the Linux and Mac OS X platforms thus allowing deployment in a small package of the FusionGrid kit with the FusionGrid login and secure MDSplus connections.

To enable fine-grained controlled resource sharing, FusionGrid worked with the Globus developers to deploy callouts from the job manager to the Akenti [19,20] distributed authorization system to authorize job execution requests [21,22] and developed a system for dynamic account creation and management of local accounts [23]. Akenti allowed a richer than required access policy and resulted in unnecessary complication when granting access permissions to users. Based on user feedback related to ease-of-use issues, a Resource Oriented Authorization Management (ROAM) system [24] was designed and implemented to handle authorization on FusionGrid in place of Akenti. ROAM follows Akenti in using the X.509 and proxy certificates to authenticate and identify users, and takes advantage of the Globus gatekeeper and job manager authorization callouts that were implemented by the Globus team to facilitate Akenti and other external authorization procedures. The ROAM system provides a mechanism for centrally managing user authorization for access to fusion grid resources and was implemented using an Apache web server and a PostgreSQL database. A web site (roam.fusiongrid.org) permits users to request, view or manage resource authorizations. The Globus servers for gridftp and gatekeeper were extended using the gridmap file callout capability to query the ROAM system to determine if a user is authorized to access the server. Likewise the MDSplus data server was extended to query ROAM for connection authorization. Fusion grid resource providers can implement complex policies for resource authorization using ROAM. For instance a resource provider can first check if the user is authorized to access the site before checking for more specific authorization. In this way site security managers can disable all or a specific user from access to any participating resources at the site. A dedicated ROAM server was configured and was mirrored at MIT and LBNL for rapid failover and with nightly backups.

2.2. FUSIONGRID DATA ACCESS

Data access on FusionGrid was made available using the MDSplus data acquisition and data management system combined with the relational database Microsoft SQL server. In the fusion community, MDSplus is used at more than 30 sites, spread over 4 continents, and is 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 loss in speed or functionality [9]. Data access over high-latency wide area networks was further accelerated by implementing a parallel I/O layer underneath MDSIP, the remote access interface for MDSplus. After initial prototyping and testing, a production version was implemented using the BIDI (bi-directional e-mode) driver under development at ANL as part of the Globus/XIO subsystem. The main MDSplus data repositories at the three experimental facilities (Alcator C-Mod [25], DIII-D [26], NSTX [27]) are securely accessible via FusionGrid. Data

management by MDSplus of large datasets generated by simulation codes is being tested at NERSC [28] using results from NIMROD simulations. NIMROD is a 3D MHD simulation code [29] that runs on large parallel computers [30], which is partially funded by the SciDAC Center for Extended MHD Modeling. The SQL servers store high-level data and metadata including the run management database and FusionGrid monitoring systems. SQL Server is securely accessible via MDSplus pending a production release of Globus on Windows.

Three prototype MDSplus web services were developed to provide some experience in the development and use of web services for fusion research. One web service accepts the name of an MDSplus server, a tree name, an experiment instance (shot number) and an expression to be evaluated. The web service returns a string representing the answer of the evaluation. A second web service was developed to accept the same parameters but returns a floating-point array for the answer. The third web service again accepts the same input parameters but returns a more complex structure called a signal consisting of a y-axis and units and an x-axis and units. The web services were developed using Microsoft Visual Studio .NET and do not provide secure communications.

2.3. FUSIONGRID COMPUTATIONAL CODES

The code TRANSP, used for time dependent analysis and simulation of tokamak plasmas, was released in late 2002 as a service on FusionGrid. Running on a Linux cluster at Princeton Plasma Physics Laboratory (PPPL), this FusionGrid service has performed over 7500 simulations taking over 50,000 CPU hours for 10 different experimental fusion devices. Both the input data and output data of TRANSP is securely written and read from an MDSplus and SQL data repository. Work is being completed to allow coordinated operation of TRANSP, a serial code, with TORIC, an RF heating code, running on a parallel cluster. The approach employed could be a model for other “hybrid” calculations. Monitoring and tracking of these runs has been accomplished through the deployment of FusionGrid’s Monitoring System (FGM) [31], which is a derivative of the Data Analysis Monitoring System used at DIII-D. FGM was designed to allow for asynchronous tracking of multiple Grid applications, and to provide logfile access through anonymous FTP. FGM is based on Java Servlet technology and is presented in an intuitive manner to the scientist via a web interface. Detailed scientific information on the progress of each run can also be obtained via FGM by examining log files linked on the website and graphically with a link to EIVis a collaborative Java-based scientific graphics application which can be run from a web browser. Thus a reasonably comprehensive web-based environment has been provided in support of production of plasma computations. This environment has evolved in response to user comments and is readily extensible to simulation applications of increasing complexity.

Deployment of TRANSP on FusionGrid frees scientists from the need to build and maintain local versions of a large and complex code, and at the same time, eases the burden on the code development team, which was previously required to provide support on a highly heterogeneous set of remote machines. Further, the Grid implementation enabled users to perform significantly more calculations, resulting in more rapid progress than prior to the deployment of FusionGrid. Scientific results using FusionGrid were presented at the 2002 through 2005 American Physical Society Division of Plasma Physics (APS/DPP) meetings [32,33] and the 2002 and 2004 IAEA Fusion Energy Conferences [34]. Additional fusion codes that have been released as computational services include GATO, ONETWO, ElFresco, FWR, and GENRAY/CQL3D.

2.4. FUSIONGRID ADVANCED COLLABORATION SERVICES

The goals of FusionGrid's advanced collaborative environment service was 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 offsite support of experimental operations, large group collaborations in a tokamak control room, simulation/experimental data analysis meetings, and shared code debugging.

Large-format shared display walls are being used to enhance the collaborative work environment of the tokamak control room (Fig. 2). Software tools and techniques have been developed to increase ease-of-use of tiled displays in these environments. These tools include software to automatically align projectors [35,36], and interfaces to assist users and administrators [37]. Improving data distribution and scalability among tiled displays is a critical issue [38–40]. As a prototype FusionGrid service, tiled display walls have been tested in a variety of usage modalities. A single tiled wall has been used for collocated group sharing and discussion. For more remote collaboration two geographically separated tiled walls with networking software to form a shared collaborative display. These tests include tiled wall usage during tokamak experimental operations. The success of these tests has resulted in C-Mod, DIII-D, and NSTX organizations purchasing large shared display systems for their respective control rooms. Installation of these systems has been completed and they are routinely used during experimental operation.

The installation of shared display walls in fusion control rooms has facilitated group discussion, but to make the collaborations more dynamic the scientists must be able to easily push content from their workstation to the display walls. Display sharing is critical so that other scientists can view the results that one scientist has produced and visualized on their local workstation. Unfortunately previously deployed collaboration tools had several

drawbacks including platform dependencies, limited sharing modes and session initiation difficulties that don't readily lend themselves to the task. To overcome these disadvantages the NFC project has developed and deployed tools that allow for display-sharing and multi-user interaction.



Fig. 2. The control rooms of NSTX (a), DIII-D (b), and C-Mod (c) with shared display walls being used to enhance collaboration. On the DIII-D display is also video from remote collaborators in Europe who were participating in that day's experiments.

The NFC project developed an application sharing tool based on a modified VNC protocol. VNC is a pixel-based protocol for desktop sharing and is inherently cross-platform compatible. The original VNC protocol only provides for sharing an entire desktop. This leads to privacy issues in a collaboration environment where a user may want to share a graph with other scientists but does not want them to see other applications such as email. To overcome these limitations, the VNC protocol was modified to allow for sharing application windows rather than the entire desktop. This new protocol was implemented in a client-server application called SharedAppVnc and supports the three common platforms, Linux, Mac OS X and Windows of the fusion community. This new collaborative tool was installed in the control rooms of the three large tokamaks (Alcator C-Mod, DIII-D, NSTX). By using SharedAppVnc the fusion scientists can easily push application displays from

their workstation up to the control room's collaborative display. They can also use SharedAppVnc to connect to remote scientists and collaboratively view applications.

To use the collaborative display wall effectively it is also important that all participants have the ability to give mouse and keyboard input. In particular users may want to reposition graphs for better viewing or highlight interesting results on the display wall. To allow for multi-user input a specialized window manager was developed for Linux. It allows up to seven scientists to be using a shared display simultaneously including interacting with different applications and menus. Each user gets their own uniquely color cursor which is easily distinguish from the others. The multi-cursor window manager has been deployed in the control rooms and allows scientists using Linux, Mac OS X or Windows workstations to interact together on the display wall. Feedback is presently being received during the C-Mod and DIII-D run campaigns with initial reports being positive.

The Access Grid is used by FusionGrid to create a service that enables group-to-group interaction and collaboration that improves the user experience beyond teleconferencing. Initial support included provisioning of Virtual Venues on the Argonne Venue Server for FusionGrid partners. These venues were also outfitted with multicast/unicast bridges to enable partners to join sessions independent of their multicast connectivity. While the Access Grid is considered a possible choice for enabling collaboration in Fusion control rooms as a whole, an extension was required from user feedback to enable collaboration with individual fusion researchers from their desktops. To this end, the NFC Project developed the Personal Interface to the Grid (PIG), a single machine workstation installation of the Access Grid software designed for usage by an individual scientist. This version of the software has seen successes in the Fusion environment and in other environments; single-machine Access Grid installations now account for a significant fraction of the AG user base. With increasing deployment of Apple Macintosh OS X computers in the Fusion environment, it became a priority for the Fusion community to have the Access Grid software available for this platform. The Access Grid is now fully supported under Mac OS X, and is in use on Mac computers in the Fusion community and in the wider Access Grid community. With the transition of FusionGrid security to use MyProxy, the AG venue client was augmented to support the retrieval of proxy credentials from FusionGrid MyProxy servers. This work enabled fusion researchers to use their FusionGrid credentials with the Access Grid software in the same manner as they use their credentials with other FusionGrid resources. These developments have increased Access Grid's capability to provide scientific research with a complex multi-site visual and collaborative experience integrated with high-end visualization environments including tiled display walls.

VoIP phones were deployed to the Alcator C-Mod control room for testing SIP-based collaboration tools. Engineering and physics operators are evaluating the new technology during Alcator C-Mod operations. A web-based map for locating and dialing the VoIP phones was developed; the map displays control room phones, which when clicked will cause the client phone to call the selected control room phone.

The SCIRun Problem Solving Environment (PSE) is 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 and 2D graphics. This capability allows for the exploration of complex simulation results and the comparison of simulation and experimental data. SCIRun on FusionGrid gives the scientist a no-license-cost visualization capability that rivals present day commercial visualization packages.

To accelerate the usage of SCIRun within the fusion community, a stand-alone application built on top of SCIRun was developed and deployed. This application, FusionViewer, allows users who are unfamiliar with SCIRun to quickly create visualizations and perform analysis of their simulation data from either the MDSplus data storage environment or from locally stored HDF5 files. More advanced tools for visualization and analysis also were created in collaboration with the SciDAC Center for Extended MHD Modeling [30]. Versions of SCIRun with the FusionViewer have been made available to fusion scientists on the Mac OS X, Linux, and other Unix based platforms and have downloaded 1163 times. SCIRun has been used with NIMROD, M3D, BOUT fusion simulation data as well as simulation data from other SciDAC application areas (e.g., Astrophysics). The subsequent visualization results — including animations — have been incorporated into invited talks at multiple APS/DPP meetings as well as peer-reviewed journal articles. As an example, SCIRun was used for the visualization and analysis of a NIMROD simulation of a disruption that occurred in a DIII-D experiment. The resulting animations and stills were presented as part of invited talks at APS/DPP meetings and the SC04 conference in addition to being highlighted in the NIH/NSF Visualization Research Challenges Report.

Web browser accessible java-based scientific graphics (ElVis) was prototyped in the NFC project and is used in FGM as mentioned above. The results and wide area collaboration features have been sufficiently promising that ElVis has been adopted by other projects including plasma diagnostic development for reflectometry on current tokamak collaborations and MSE for ITER. ElVis has also been adopted in future plans for large-scale simulation projects for SciDAC as an interactive graphical monitoring tool for running simulations.

2.5. COLLABORATIVE CONTROL ROOM

The capabilities of FusionGrid were unified in a demonstration at SC03 of a collaborative tokamak control room. Magnetic fusion experiments operate in a pulsed mode producing plasmas of up to 10 seconds in duration every 10 to 20 minutes. Decisions to change the next pulse are informed by data analysis conducted within the roughly 20-minute between-pulse interval. This mode of operation requires rapid data analysis that can be assimilated in near-realtime by a geographically dispersed research team. Offsite collaborators (Phoenix) joined in a mockup of a DIII-D experiment located in San Diego. AG technology allowed for shared audio and video as well as shared applications. The offsite collaborators could hear DIII-D announcements and see via a web interface the state of the pulse cycle, the status of data acquisition, and the state of between pulse data analysis. The fusion visualization application ReviewPlus was shared between the two sites allowing for joint scientific exploration. Between pulse data analysis of the plasma shape (EFIT running at PPPL) was conducted on FusionGrid through a computational reservation system that guaranteed a specific analysis to be completed within a set time window [41]. Additionally, the TRANSP computational service was run at PPPL for the first time between pulses giving scientists data that was previously only available after the experimental day had ended. The offsite team was able to collaborate more efficiently amongst themselves by being able to share their personal display with the room's shared display. This capability allowed visualizations to be efficiently compared side-to-side for debate before reporting results back to the DIII-D control room.

Based on the positive SC03 demonstration, in January 2004, a scientist from San Diego decided not to travel and instead remotely lead the experiment at the JET tokamak facility in England. This successful remote leadership was made possible by the same technology demonstrated at SC03. Today's experiments invariably involve a mix of local and remote researchers, with sessions routinely led by scientists at remote institutions. For example, over the last 2 years, 5-10% of experimental time on the Alcator C-Mod tokamak was devoted to runs with remote session leaders. From our experience of placing collaborative tools in operating tokamak control rooms, the project has determined that the collaborative control room requires (1) secured computational services that can be scheduled as required, (2) the ability to rapidly compare experimental data with simulation results, (3) a means to easily share individual results with the group by moving application windows to a shared display, and (4) the ability for remote scientists to be fully engaged in experimental operations through shared audio, video, and applications.

In an effort to assess the ability to schedule secured computational services, the NFC project conducted a successful advanced reservation test on FusionGrid that was operated in parallel with DIII-D experimental operations. Since that first test the advance reservation

system has been improved by building more semantically rich and user-friendly interfaces as well as by examining new technologies to create more reliable remote computation access for experiments. Specifically, virtual machines (VM) were examined based on concepts [42,43] coupled with advance reservation systems to guarantee and enforce specified resource quota during experiments. By providing live migration capabilities VMs allow for much more flexible and failure-independent management of remote resources at the cost of very little overhead thus providing a promising platform for collaborative on-demand resource sharing.

3. OVERALL ASSESSMENT

The project would be considered a success if assessed by its initial objective to develop and deploy a national FES “Grid” (FusionGrid) that would be a system for secure sharing of computation, visualization, and data resources over the Internet. But even more important, FusionGrid has had and continues to have a tremendous impact on how collaboration and computational services are delivered to the fusion community — specially the experimental community and also those users who are not themselves authors of services being utilized. What started out as a prototype or pilot project, FusionGrid services proved so successful they are now used in a production environment for everyday FES research. The success of other prototype services has resulted in FES research funds being used to purchase necessary hardware to create production services. As a result, FusionGrid has had a substantial impact on how the community performs its science.

The project has also been assessed from the outside including annually by the DOE/OFES Program Advisory Committee (PAC) of the Plasma Science Advanced Computing Institute (PSACI). Chaired by Dr. W.L. Kruer, the PAC is comprised of leading experts in plasma physics and computer science. Quoting from their summer of 2005 report:

“The National Fusion Collaboratory (NFC) has continued to be remarkably successful in providing vital infrastructure for national and international collaboration in experimental operation, data analysis, and simulation. FusionGrid services have enabled fusion scientists to use codes with less effort and better support. The code TRANSP, available as a FusionGrid service, has performed over 5800 simulations for ten different fusion devices throughout the world and is also being used for machine and experimental designs for both the Burning Plasma Experiment in the US and the International Thermonuclear Experimental Reactor. Recently the GATO ideal MHD plasma stability code was also released as a FusionGrid service. Shared display walls in Tokamak control rooms and significantly enhanced remote collaboration via Access Grid nodes are having a large impact on productivity. The NFC is enabling greater utilization of the existing US experiments by more scientists by facilitating real-time off-site interaction with the experiments. In addition, a unified FusionGrid security infrastructure has enabled more effective sharing of data and codes.”

Given this positive review they recommended in the same report that:

“...the remarkably successful National Fusion Collaboratory continue to play a major leadership role in building collaborative infrastructure for future remote operation of experiments and in fostering collaborative development and use of simulation codes as well as data analysis codes. The

NFC is encouraged to take a leadership role in establishing common data structures and protocols and a collaborative culture for integrating SciDAC and FSP project results into the broad fusion community, and in particular the large experimental communities.”

Looking toward the future, the fusion program will be increasingly oriented toward ITER where even now, a decade before operation begins, a large portion of national program efforts are organized around coordinated efforts to develop promising operational scenarios. Substantial efforts to develop integrated plasma modeling codes are also underway in the U.S., Europe and Japan. While the NFC project has helped to accelerate the development and exploitation of remote collaboration technology and helped to place the U.S. in a leadership position with respect to enabling technology for international collaborations, with such a large investment at stake for the US fusion program, there is a clear need to improve the integration and reach of available tools.

ACKNOWLEDGMENT

Work supported by the U.S. Department of Energy under DE-FC02-01ER25455, DE-AC02-76CH03073, DE-FG02-04ER54235, W-31-109-ENG-38, DE-FC02-01ER25454 and DE-AC03-76SF00098. The authors gratefully acknowledge the valuable conversations and feedback from the many fusion scientists at Alcator C-Mod, DIII-D, and NSTX. The project gratefully acknowledges the cooperation of Jim Basney at NSCA for integrating our extensions of the MyProxy software back into the released product.

REFERENCES

- [1] D.P. Schissel for the Collaboratory Team, “Building the U.S. National Fusion Grid: Results from the National Fusion Collaboratory Project,” *Fusion Engin. Design* **71**, 245–250 (2004).
- [2] <http://www.fusiongrid.org/projects/>
- [3] <http://www.scidac.org/>
- [4] S. Dean, “Fifty years of U.S. Fusion Research — An Overview of Programs,” *Nuclear News* (American Nuclear Society, 2002) p. 34.
- [5] F.W. Perkins, “ITER Physics Basis: Overview and Summary.” *Nucl. Fusion* **39**, 2137–74 (1999).
- [6] <http://www.wofe.er.doe.gov/>
- [7] T. Fredian, J. Stillerman, “MDSplus Remote Collaboration Support-Internet and World Wide Web,” *Fusion Engin. Design* **43**, 327 (1999).
- [8] <http://www.mdsplus.org/>
- [9] I. Foster, C. Kesselman, G. Tsudik, and S. Tuecke, “A Security Architecture for Computational Grids,” *Proc. 5th ACM Conf. on Computer and Communications Security* (1998).
- [10] K. Keahey, T. Fredian, et al., “Computational Grids in Action: the National Fusion Collaboratory,” *Future Generation Computer Systems, The International Journal of Grid Computing: Theory, Methods and Applications* **18**(8), 1005–1015 (2002).
- [11] R.J. Hawryluk, “An Empirical Approach to Tokamak Transport,” *Phys. Plasmas Close to Thermonuclear Conditions*, ed. by B. Coppi, et al., (CEC, Brussels, 1980) Vol. 1, p. 19.
- [12] J. Burruss, et al., “Remote Computing on the National Fusion Grid,” 4th IAEA Tech. Mtg. on Control, Data Acquisition, and Remote Participation for Fusion Research, July 21–23, 2003, San Diego, California to be published in *Fusion Engineering and Design*.
- [13] <http://www.accessgrid.org/>
- [14] L. Childers, T. Disz, R. Olson, M.E. Papka, R. Stevens, and T. Udeshi, “Access Grid: Immersive Group-to-Group Collaborative Visualization,” *Proc. 4th Intl. Immersive Projection Technology Workshop*, Ames, Iowa, 2000.
- [15] Scalable Display Wall Project, Princeton University, <http://www.cs.princeton.edu/omnimedia/>
- [16] S.G. Parker, D.M. Weinstein and C.R. Johnson, “The SCIRun Computational Steering Software System,” *Modern Software Tools in Scientific Computing* (Birkhauser Press, 1997) pp. 1–44.

- [17] J.R. Burruss, T.W. Fredian, M.R. Thompson, “Simplifying FusionGrid Security,” Challenges of Large Applications in Distributed Environments (CLADE) Workshop at HPDC14, July 2005, Research Triangle Park, NC; LBNL-58691.
- [18] <http://www.pubcookie.org/>
- [19] M. Thompson, A. Essiari, S. Mudumbai, “Certificate-Based Authorization Policy in a PKI Environment,” ACM Transactions on Information and System Security (TISSEC) **6.4**, 566–588 (2003).
- [20] M. Thompson, A. Essiari, K. Keahey, V. Welch, S. Lang, B. Liu, “Fine-grained Authorization for Job and Resource Management using Akenti and Globus,” 2003 Conf. for Computing in High Energy and Nuclear Physics, La Jolla, California, 2003.
- [21] K. Keahey, V. Welch, et al. “Fine-Grained Authorization for Job Execution in the Grid: Design and Implementation,” Concurrency and Computation: Practice and Experience **16**(5), 477-488 (2004).
- [22] K. Keahey, and V. Welch, “Fine-Grain Authorization for Resource Management in the Grid Environment,” to be published in Proc. Grid 2002 Workshop (November 2002).
- [23] K. Keahey, M. Ripeanu, et al. “Dynamic Creation and Management of Runtime Environments in the Grid,” Workshop on Designing and Building Web Services, Chicago, Illinois, 2003.
- [24] J.R. Burruss, T.W. Fredian, M.R. Thompson, “Security on the US Fusion Grid,” Fusion Engin. Design **81**(15-17) 1949-1955 (2006).
- [25] I. Hutchinson, R. Boivin, F. Bombarda, et al., “Alcator C-Mod,” Phys. Plasmas **1**, 1511 (1994).
- [26] J.L. Luxon and L. Davis, Fusion Technol. **8**, 441 (1985).
- [27] S. Kaye, et al., Fusion Technol. **36**, 16 (1999).
- [28] National Energy Research Scientific Computing Center (NERSC), <http://www.nersc.gov/>.
- [29] A.H. Glasser, C.R. Sovinec, R.A. Nebel, T.A. Gianakon, S.J. Plimpton, M.S. Chu, D.D. Schnack, and the NIMROD Team, “The NIMROD Code: A New Approach to Numerical Plasma Physics,” Plasma Phys. Control. Fusion **41**, A747 (1999).
- [30] <http://w3.pppl.gov/CEMM/>
- [31] S. Flanagan, et al., “A General Purpose Data Analysis Monitoring System with Case Studies from the National Fusion Grid and the DIII-D MDSplus between Pulse Analysis System,” Fusion Engin. Design **71**, 263-267 (2004).
- [32] D. Brennan, et al., Phys. Plasmas **10**, 1643 (2003).
- [33] M. Murakami, et al., Phys. Plasmas **10**, 1691 (2003).

- [34] E.J. Doyle, et al., “Core and Edge Aspects of Quiescent Double Barrier Operation on DIII-D with Relevance to Critical ITB Physics Issues,” Proc. 19th IAEA Fusion Energy Conf., Lyon, France, 2002, Paper EX/C3-2; available online at <http://www.iaea.org/programmes/ripc/physics/fec2002/html/fec2002.htm>.
- [35] G. Wallace, H. Chen, and K. Li, “DeskAlign: Automatically Aligning a Tiled Windows Desktop,” IEEE Intl. Workshop on Projector-Camera Systems (PROCAMS2003), October 2003.
- [36] H. Chen, R. Sukthankar, G. Wallace, and K. Li, “Scalable Alignment of Large-Format Multi-Projector Displays Using Camera Homography Trees,” Proc. Conf. on Visualization '02 Session P11 (2002) pp. 339–346.
- [37] G. Wallace, “Display Wall Cluster Management,” IEEE Visualization 2002 Workshop on Commodity-Based Visualization Clusters.
- [38] H. Chen, G. Wallace, A. Gupta, K. Li, T. Funkhouser, P. Cook, “Experiences with Scalability of Display Walls,” Tom Funkhouser, Perry Cook. to appear in Immersive Projection Technology Symposium (IPT), March 2002.
- [39] Han Chen, Yuqun Chen, Adam Finkelstein, Thomas Funkhouser, Kai Li, Zhiyan Liu, Rudrajit Samanta, and Grant Wallace, “Data Distribution Strategies for High-Resolution Displays,” Computers & Graphics, Special Issue on Mixed Realities — Beyond Conventions **25**, 811 (2001).
- [40] Yuqun Chen, Han Chen, Douglas W. Clark, Zhiyan Liu, Grant Wallace, and Kai Li, “Software Environments for Cluster-based Display Systems,” The First IEEE/ACM Intl. Symp. on Cluster Computing and the Grid (CCGrid 2001), Brisbane, Australia, 2001.
- [41] K. Keahey, and K. Motawi, “Taming of the Grid: Virtual Application Services,” Argonne National Laboratory, Mathematics and Computer Science Division Technical Memorandum ANL/MCS-TM-262 (2003).
- [42] Virtual Workspaces: <http://workspace.globus.org>
- [43] J. Burruss and K. Keahey, “The Flexible Execution (FLEX) Cluster,” NFC White Paper, January, 2006.

APPENDIX PUBLICATIONS LIST

- Abla, G., et al., "Shared Display Wall Based Collaboration Environment in the Control Room of the DIII-D National Fusion Facility," Proc. Wkshp. on Advanced Collaborative Environments, Redmond, Washington, 2005. (<http://www.mcs.anl.gov/fl/fevents/wace/wace2005/>)
- Abla, G., Flanagan, S.M., Peng, Q., Burruss, J.R., Schissel, D.P., "Advanced Tools for Enhancing Control Room Collaborations," Proc. 5th IAEA Tech. Mtg. on Control, Data Acquisition, and Remote Participation for Fusion Research, July 12-15, 2005, Budapest, Hungary; *Fusion Engin. Design* **81**(15-17), 2039-2044 (2006).
- Allcock, W., Bester, J., Bresnahan, J., Foster, I., Gawor, J., Insley, J., Link, J., Papka, M., "GridMapper: A Tool for Visualizing the Behavior of Large-Scale Distributed Systems," Proc. 11th IEEE Intl. Symp. on High Performance Distributed Computing (2002) pp. 179-188.
- Binns, J., Papka, M., Stevens, R., Peng, Q., Schissel, D., "Development of Cluster-based Image Viewer," submitted to 11th IEEE Intl. Symp. on High Performance Distributed Computing, July 24-26, 2002, Edinburgh, Scotland.
- Burruss, J.R., "The Resource Oriented Authorization Manager (ROAM)," Proc. 14th IEEE Intl. Symp. on High Performance Distributed Computing (2005) pp. 308-309.
- Burruss, J.R., et al., "Developments in Remote Collaboration and Computation," *Fusion Sci. Technol.* **47**(3), 814 (2005).
- Burruss, J.R., et al., "Remote Computing on the National Fusion Grid," *Fusion Engin. Design* **71**, 251-255 (2004)
- Burruss, J.R., et al., "ROAM: An Authorization Manager for Grids," accepted for publication in *J. Grid Computing*, 2006.
- Burruss, J.R., Fredian, T.W., Thompson, M.R., "Security on the U.S. Fusion Grid," 5th IAEA Tech. Committee Mtg. on Control, Data Acquisition and Remote Participation for Fusion Research, July 12–15, 2005, Budapest, Hungary; *Fusion Engin. Design* **81**(15-17), 1949-1955 (2006).
- Burruss, J.R., Fredian, T.W., Thompson, M.R., "Simplifying FusionGrid Security," Proc. Challenges of Large Applications in Distributed Environments (2005) pp. 95-105.
- Chen, H., et al., "A Parallel Ultra-High Resolution MPEG-2 Video Decoder for PC Cluster Based Tiled Display System," Proc. Intl. Parallel and Distributed Processing Symp., 2002.
- Chen, H., et al., "Calibrating Scalable Multi-Projector Displays Using Camera Homography Trees," Proc. Computer Vision and Pattern Recognition, 2001.

- Chen, H., et al., “Data Distribution Strategies for High-Resolution Displays,” *Computers and Graphics* **25**(5), 811-818 (2001).
- Chen, H., et al., “Experiences with Scalability of Display Walls,” Proc. 7th Annual Immersive Projection Technology Symp., 2002.
- Chen, H., et al., “Memory Performance Optimizations for Real-Time Software HDTV Decoding,” *J. VLSI Signal Processing* **41**(2), 193-207 (2005).
- Chen, H., et al., “Memory Performance Optimizations for Real-Time Software HDTV Decoding,” Proc. IEEE Intl. Conf. on Multimedia and Expo, 2002.
- Chen, H., et al., “Scalable Alignment of Large-Format Multi-Projector Displays Using Camera Homography Trees,” Proc. IEEE Visualization, 2002.
- Chen, Y., et al., “Software Environments for Cluster-based Display Systems,” Proc. 1st IEEE/ACM Intl. Symp. on Cluster Computing and the Grid, Denver, Colorado, 2001.
- Feibush, E., Kramer, G., Valeo, E., Nazikian, R., McCune, D., “Multi-tier Graphical Web Service for Simulating Reflectometry in Plasma,” ACM SIGGRAPH 2006 Conf. Abstracts and Applications, July 2006, p. 420
- Flanagan, S., et al., “A General Purpose Data Analysis Monitoring System with Case Studies from the National Fusion Grid and the DIII-D MDSplus between Pulse Analysis System,” *Fusion Engin. Design* **71**, 263-267 (2004).
- Foster, I., “Service-Oriented Science,” *Science Magazine* **308**(5723), 814-817 (2005).
- Foster, I., Freeman, T., Keahey, K., Scheftner, D., Sotomayor, B., Zhang, X., “Virtual Clusters for Grid Communities,” Proc. 6th IEEE Intl. Symp. on cluster Computing and the Grid, May 16-19, 2006, Singapore.
- Fredian, T.W., Stillerman, J.A., “MDSplus, Current Developments and Future Directions,” *Fusion Engin. Design* **60**, 229-233 (2002).
- Greenfield, C.M., “Advanced Tokamak Research at the DIII-D National Fusion Facility in Support of ITER,” *Journal of Physics: Conference Series* **16**, 586-590 (2005).
- Greenwald, M., Schissel, D.P., Burruss, J.R., Fredian, T., Lister, J., Stillerman, J., “Visions for Data Management and Remote Collaboration on ITER,” Proc. 10th Intl. Conf. on Accelerator and Large Experimental Physics Control Systems, Geneva, Switzerland, 2005.
- Keahey, K., “Fine Grain Authorization Policies in the Grid: Design and Implementation,” Proc. ACM/IFIP/USENIX Intl. Middleware Conf., June 16-20, 2003, Rio de Janeiro, Brazil.
- Keahey, K., et al., “Agreement Based Interactions for Experimental Science,” Proc. Euro-Par 2004, August 31-September 3, 2004, Pisa, Italy.

- Keahey, K., et al., “Computational Grids in Action: The National Fusion Collaboratory,” *Future Generation Computing Systems* **18**(8), 1005-1015 (2001).
- Keahey, K., et al., “Fine-Grained Authorization for Job Execution in the Grid: Design and Implementation,” *Concurrency and Computation: Practice and Experience* **16**(5), 477-488 (2004).
- Keahey, K., et al., “Grids for Experimental Science: The Virtual Control Room,” Proc. 2nd Intl. Workshop on Challenges of Large Applications in Distributed Environments, (2004) p. 4.
- Keahey, K., Foster, I., Freeman. T., Zhang, X., Galron, D., “Virtual Workspaces in the Grid,” Proc. Euro-Par 2005, August 30-September 2, 2005, Lisbon, Portugal.
- Keahey, K., Motawi, K., “Taming the Grid: Virtual Application Services,” ANL Technical Memorandum ANL/MCS-TM-262, 2003.
- Keahey, K., Welch, V., “Fine-Grain Authorization for Resource Management in the Grid Environment,” *Lecture Notes in Computer Science*, Grid Computing—GRID 2002: Third International Workshop, Baltimore, MD, USA, November 18, 2002. Proceedings, M. Parashar (Ed.) (Springer, 2002) vol. 2536, p. 199.
- Kruger, S.E., et al., “Dynamics of the Major Disruption of a DIII-D Plasma,” *Phys. Plasmas* **12**, 056113 (2005).
- Li, K., et al., “Dynamic Scalable Visualization for Collaborative Scientific Applications,” Proc. Next Generation Software Workshop, April 4, 2005, Denver, Colorado.
- Liu, Z., et al., “Improving Progressive View-Dependent Isosurface Propagation,” *Computers and Graphics* **26**(2), 209-218 (2002).
- Liu, Z., et al., “Progressive View-Dependent Isosurface Propagation,” Proc. VisSym ’01, Joint Eurographics-IEEE TCVG Symp. Visualization, May 28-30, 2001, Ascona, Switzerland.
- Minor, D.H., Burruss, J.R., “Distributed MDSplus Database Performance with Linux Clusters,” 5th IAEA Tech. Committee Mtg. on Control, Data Acquisition and Remote Participation for Fusion Research, July 12–15, 2005, Budapest, Hungary; *Fusion Engin. Design* **81**(15-17), 1999-2002 (2006).
- Samanta, R., et al., “Parallel Rendering with K-Way Replication,” Proc. IEEE Symp. on Parallel and Large-Data Visualization and Graphics, 2001.
- Sanderson, A.R., et al., “Display of Vector Fields Using a Reaction-Diffusion Model,” Proc. 15th IEEE Visualization 2004 Conf. (VIS 2004), October 10-15, 2004, Austin, Texas, USA (IEEE Computer Society 2004, ISBN 0-7803-8788-0) pp. 115-122.
- Sanderson, A.R., Kirby, M., Johnson, C.R., Yang, L., “Advanced Reaction-Diffusion Models for Texture Synthesis,” accepted for publication in *J. Graphics Tools*, 2006.

- Sanderson, A.R., Meyer, M.D., Kirby, R.M., Johnson, C.R., “A Framework for Exploring Numerical Solutions of Advection-Reaction-Diffusion Equations Using a GPU-Based Approach,” submitted to Computing and Visualization in Science, 2006.
- Schissel, D.P., “An Advanced Collaborative Environment to Enhance Magnetic Fusion Research,” Proc. 2nd Wkshp. on Advanced Collaborative Environments (WACE) held at the 2nd IEEE Intl. Symp. on Cluster Computing and the Grid (CCGrid 2002), May 21-24, 2002, Berlin, Germany.
- Schissel, D.P., “Advances in Remote Participation for Fusion Experiments,” to be published in Proc. 21st IEEE/NPSS Symp. on Fusion Engineering, September 26-29, 2005, Knoxville, Tennessee.
- Schissel, D.P., “Advances in the Real-Time Interpretation of Fusion Experiments,” Journal of Physics: Conference Series **16**, 49-53 (2005).
- Schissel, D.P., “National Fusion Collaboratory — Advancing the Science of Fusion Research,” Proc. Proc. 2003 DOE SciDAC PI Meeting March 10-11, 2003, Napa, California.
- Schissel, D.P., “The National Fusion Collaboratory Project,” Proc. of 2004 DOE SciDAC PI Meeting, March 22-24, 2004, Charleston, North Carolina.
- Schissel, D.P., et al., “Advances in Remote Participation for Fusion Experiments,” Fusion Engin. Design **74**, 803 (2005).
- Schissel, D.P., et al., “Building the U.S. National Fusion Grid: Results from the National Fusion Collaboratory Project,” Fusion Engin. Design **71**, 245-250 (2004).
- Schissel, D.P., et al., “Data Management, Code Deployment, and Scientific Visualization to Enhance Scientific Discovery in Fusion Research through Advanced Computing,” Fusion Engin. Design **60**, 481-486 (2002).
- Schissel, D.P., et al., “Grid Computing and Collaboration Technology in Support of Fusion Energy Sciences,” Phys. Plasmas **12**, 058104 (2005).
- Schissel, D.P., et al., “The National Fusion Collaboratory Project: Applying Grid Technology for Magnetic Fusion Research,” Proc. Workshop on Case Studies on Grid Applications in conjunction with GGF10, March 13, 2004, Berlin, Germany.
- Schissel, D.P., for the National Fusion Collaboratory Project, “The Collaborative Tokamak Control Room,” 5th IAEA Tech. Committee Mtg. on Control, Data Acquisition and Remote Participation for Fusion Research, July 12–15, 2005, Budapest, Hungary; Fusion Engin. Design **81**(15-17), 2031-2037 (2006).
- Schissel, D.P., Greenwald, M.J., Johnston, W.E., “Fusion Energy Sciences Networking: A Discussion of Future Requirements,” Proc. DOE High Performance Network Planning Workshop, August 13-15, 2002, Reston, Virginia.

- Simpson, J., Sanderson, A., Luke, E., Balling, K., Coffey, R., “Collaborative Remote Visualization,” Proc. IEEE/ACM SC2002 Conf. on High Performance Networking and Computing, November 16-22, 2002, Baltimore, Maryland.
- Thompson, M., “Fine-grained Authorization for Job and Resource Management using Akenti and Globus,” Proc. 2003 Conf. for Computing in High Energy and Nuclear Physics, March 24-28, 2003, La Jolla, California.
- Wallace, G., “Display Wall Cluster Management,” Proc. IEEE Visualization 2002 Workshop on Commodity-Based Visualization Clusters, October 27–November 1, 2002, Boston, Massachusetts.
- Wallace, G., et al., “Automatic Alignment of Tiled Displays for a Desktop Environment,” *J. Software* **15**(12), 1776-1786 (2004).
- Wallace, G., et al., “Color Gamut Matching for Tiled Display Walls,” Proc. 7th Intl. Wkshp. on Immersive Projection Technology, 9th Eurographics Wkshp. on Virtual Environments, 2003, pp. 293-302
- Wallace, G., et al., “DeskAlign: Automatically Aligning a Tiled Windows Desktop,” Proc. IEEE Intl. Workshop on Projector-Camera Systems, October 12, 2003, Nice, France.
- Wallace, G., et al., “A MultiCursor X Window Manager Supporting Control Room Collaboration,” Princeton University, Computer Science, Technical Report TR-707-05, (2004).
- Wallace, G., et al., “Tools and Applications for Large-Scale Display Walls,” *IEEE Computer Graphics & Applications*, **25**(4), 24-33 (2005).
- Zhang, H., et al., “Providing Data Transfer with QoS as Agreement-Based Service,” Proc. 2004 IEEE Intl. Conf. on Services Computing, September 15-18, 2004, Shanghai, China.