LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

UCRL-JRNL-232762

Building a Global Federation System for Climate Change Research: The Earth System Grid Center for Enabling Technologies (ESG-CET)

R. Ananthakrishnan, D. E. Bernholdt, S. Bharathi, D. Brown, M. Chen, A. L. Chervenak, L. Cinquini, R. Drach, I .T. Foster, P. Fox, D. Fraser, K. Halliday, S. Hankin, P. Jones, C. Kesselman, J. D. E. Middleton, J. Schwidder, R. Schweitzer, R. Schuler, A. Shoshani, F. Siebenlist, A. Sim, W. G. Strand, N. Wilhelmi, M. Su, D. N. Williams

July 16, 2007

Journal of Physics: Conference Series, SciDAC #07 conference proceedings

Disclaimer

Building a Global Federation System for Climate Change Research: The Earth System Grid Center for Enabling Technologies (ESG-CET)

The Earth System Grid Center for Enabling Technologies Team:

R Ananthakrishnan[1], D E Bernholdt[7,9], S Bharathi[8], D Brown[5], M Chen[7], A L Chervenak[8], L Cinquini[5], R Drach[3], I T Foster[1,9], P Fox[5], D Fraser[1], K Halliday[3], S Hankin[6], P Jones[4], C Kesselman[8], J D E Middleton[5,9], J Schwidder[7], R Schweitzer[6], R Schuler[8], A Shoshani[2], F Siebenlist[1], A Sim[2], W G Strand[5], N. Wilhelmi[5], M Su[8], D N Williams[3,9]

[1] Argonne National Laboratory, Chicago, IL, USA.
[2] Lawrence Berkeley National Laboratory, Berkeley, CA, USA.
[3] Lawrence Livermore National Laboratory, Livermore, CA, USA.
[4] Los Alamos National Laboratory, Los Alamos, NM, USA.
[5] National Center for Atmospheric Research, Boulder, CO, USA.
[6] National Oceanic and Atmospheric Administration (PMEL), Seattle, WA, USA.
[7] Oak Ridge National Laboratory, Oak Ridge, TN, USA.
[8] University of Southern California, Information Sciences Institute, Marina del Ray, CA, USA.
[9] E-mail: williams13@llnl.gov, don@ucar.edu, itf@mcs.anl.gov, bernholdtde@ornl.gov

Abstract. The recent release of the Intergovernmental Panel on Climate Change (IPCC) 4th Assessment Report (AR4) has generated significant media attention. Much has been said about the U.S. role in this report, which included significant support from the Department of Energy through the Scientific Discovery through Advanced Computing (SciDAC) and other Department of Energy (DOE) programs for climate model development and the production execution of simulations. The SciDAC-supported Earth System Grid Center for Enabling Technologies (ESG-CET) also played a major role in the IPCC AR4: all of the simulation data that went into the report was made available to climate scientists worldwide exclusively via the ESG-CET. At the same time as the IPCC AR4 database was being developed, the National Center for Atmospheric Research (NCAR), a leading U.S. climate science laboratory and a ESG participant, began publishing model runs from the Community Climate System Model (CCSM), and its predecessor the Parallel Coupled Model (PCM) through ESG. In aggregate, ESG-CET provides seamless access to over 250 terabytes of distributed climate simulation data to over 6,000 registered users worldwide, who have taken delivery of more than 280 terabytes from the archive. Not only does this represent a substantial advance in scientific knowledge, it is also a major step forward in how we conduct the research process on a global scale. Moving forward, the next IPCC assessment report, AR5, will demand multi-site metadata federation for data discovery and cross-domain identity management for single sign-on of users in a more diverse federation enterprise environment. Towards this aim, ESG is leading the effort in the climate community towards standardization of material for the global federation of metadata, security, and data services required to standardize, analyze, and access data worldwide.

1. Introduction

The U.S. Department of Energy's (DOE) investment in climate change research is broad, spanning model development, climate change simulation, model intercomparisons, observational programs, and supporting infrastructure for the Intergovernmental Panel on Climate Change (IPCC) [2]. Increasingly, climate change research is data intensive, involving the analysis and intercomparison of simulation and observation data from many sources. Continued scientific progress depends upon powerful, effective enabling technologies that allow the core climate science community to coherently manage and publish a diverse collection of what in a few years will be petascale data, such that a broad, global community can access and analyze it. The ability of analysts to integrate data from a variety of sources (multiple types of models and observations, as well as economic, agricultural, and other data) is emerging as another important direction in climate change research.

DOE's SciDAC program brought major advances in not only climate modeling, via the Community Climate System Model (CCSM) [8] effort, but—equally importantly—in the management and sharing of distributed terascale data, via the Earth System Grid (ESG) project [5, 9]. The multidisciplinary ESG team developed and delivered a production environment that serves a worldwide community with climate data from multiple climate model sources (e.g., CCSM, PCM, IPCC), ocean model data (e.g., POP), CCSM source code distribution, and analysis and visualization tools. Data holdings are distributed across multiple sites including LANL, LBNL, LLNL, NCAR, and ORNL. ESG also operates a dedicated portal that supports the IPCC community in the development of its 4th Assessment Report (AR4) [2]. ESG now supports over 6,000 registered users from around the globe; manages over 250 TB of data, models, and tools; and has delivered more than 280 TB of data to its users. It is estimated that over 300 scientific publications [3] have been authored focused upon the analysis of the IPCC data alone.

Development continues on the Earth System Grid environment, responding both to the experience gained through the publication of CCSM, IPCC AR4, and other major holdings as well as to the future needs of the community for increasing levels of federation with other modeling archives and other types of environmental data around the world.

2. A User's View of ESG

ESG is a large, production, distributed system with primary access currently via two web portals: one for general climate research data and another dedicated to the IPCC activity.

Users of the ESG portal first undergo a registration process [10], where they are made known to the system and granted various privileges and access to data collections. Group-based authorization mechanisms allow control over which users can access which data, a common requirement of data "owners" or curators in this field.

The main portal page, shown in Figure 1, provides news, status, and live monitoring of ESG. Once logged in, users may either search or browse ESG catalogs to locate desired datasets, with the option of browsing both collection-level and file/usage-level metadata. Based on this perusal of the catalogs, users may gather a collection of files into a "DataCart" or request an "aggregation," which allows them to request a specific set of variables subject to a spatiotemporal constraint. Once the user's request is finalized, the aggregation is produced, or the data is retrieved from disk or deep storage at the appropriate archive site(s), and the user is notified that their data are ready to be retrieved. Users can download their data using a web browser, wget, or a DataMover Lite (DML) client, which utilizes GridFTP [4] for the actual transfer, and offers additional features to facilitate bulk movement of data. The fact that the IPCC AR4 data is all held in a single location and uniformly organized has allowed that portal to also present a direct FTP-based interface to browse and retrieve the data. Metadata remains accessible only through the web interface. Other technologies used in the ESG implementation include RLS [6], SRM [12], and OPeNDAP [11].

Several times in the course of the project, we have sought feedback from the ESG user community as to aspects of the system that work and do not work for their usage patterns. These surveys, together with more intensive one-on-one sessions between ESG interface developers and experienced climate scientists, have provided valuable input as we extend and improve ESG to meet the growing needs of

the community. This feedback has been invaluable as we design and implement the next generation of ESG, both in terms of user interface and functionality.

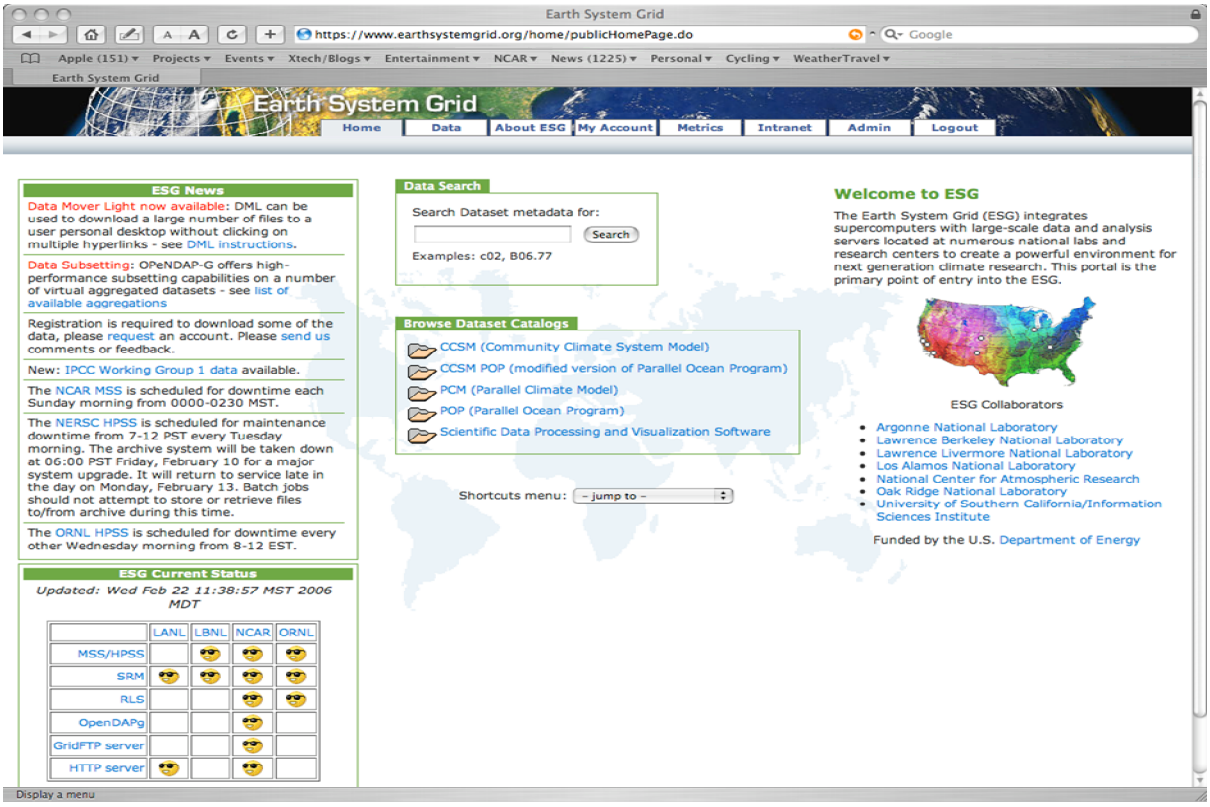

Figure 1: ESG portal, showing search form, ESG status (lower left) [7], recent news, etc.

3. Towards a Global, Federated Climate Data Community

Our high-level goals are driven primarily by scientific objectives relevant to DOE scientific priorities over the next five years, focused on three primary climate-modeling activities:

- The IPCC 5th Assessment Report (AR5) in 2010;
- The Climate Science Computational End Station (CCES) project at ORNL; and
- The SciDAC-2 application project, A Scalable and Extensible Earth System Model for Climate Change Science, which continues to develop CCSM.

The IPCC AR5 activity in particular, is a major driver for global federation. An early testbed is expected to include British Atmosphere Data Centre (U.K.), Geophysical Fluid Dynamics Laboratory (U.S.), LLNL (U.S.), Max Planck Institute for Meteorology (Germany), NCAR (U.S.), ORNL (U.S.), and the University of Tokyo Center for Climate System Research (Japan). In production, the system is expected to include upwards of 35 sites around the globe (compared to five sites now, all in the U.S.) running a combination of ESG and locally-developed software.

In addition to these primary drivers, and additional DOE-related collections already published through ESG (including the IPCC AR4, CCSM, Parallel Climate Model, and Parallel Ocean Program archives), we are also collaborating with the larger climate data community to evaluate, refine, and integrate emerging technology to establish additional national and international climate change archives, including the North American Regional Climate Change Assessment Program (NARCCAP), and the Cloud Feedback Model Intercomparison Project (CFMIP) archives. For working interoperability, standardization, and community open access issues with other data repositories, ESG is leading efforts in the Global Organization for Earth System Science Portals (GO-ESSP).

The requirements of these various projects drive not only our need for a globally federated system, but also other significant changes to the ESG architecture. For example, CCSM is adding atmospheric chemistry and land and ocean biogeochemistry to support fully coupled carbon cycle simulations and more realistic simulation of aerosol feedbacks on the climate. These new capabilities will result in substantially larger datasets, expansion of the ESG's cataloging and metadata standards, and the development and optimization of mechanisms to extract such data from much larger output files.
Larger dataset sizes have also motivated a general change in strategy. While the first generation of ESG was focused, in large part, on delivering data as published, the next generation must be more pro-active in providing users with *exactly* the data they need. Therefore, server-side processing, such as aggregation, subsetting, grid conversion, compression, and, eventually, execution of complex analysis steps, must be incorporated into the system. In doing so, however, we must also provide users of ESG server-side processing capabilities with a carefully documented provenance for any analysis results or derived data delivered to them by the system.

4. An Architecture for a Global Federation System

In order to meet the needs of the next generation ESG, the architecture must be generalized to allow a larger number of distributed sites with different capabilities to selectively federate, cooperate, or operate in standalone fashion with a variety of user access, including multiple portals and service- or API-based access, and data delivery mechanisms.

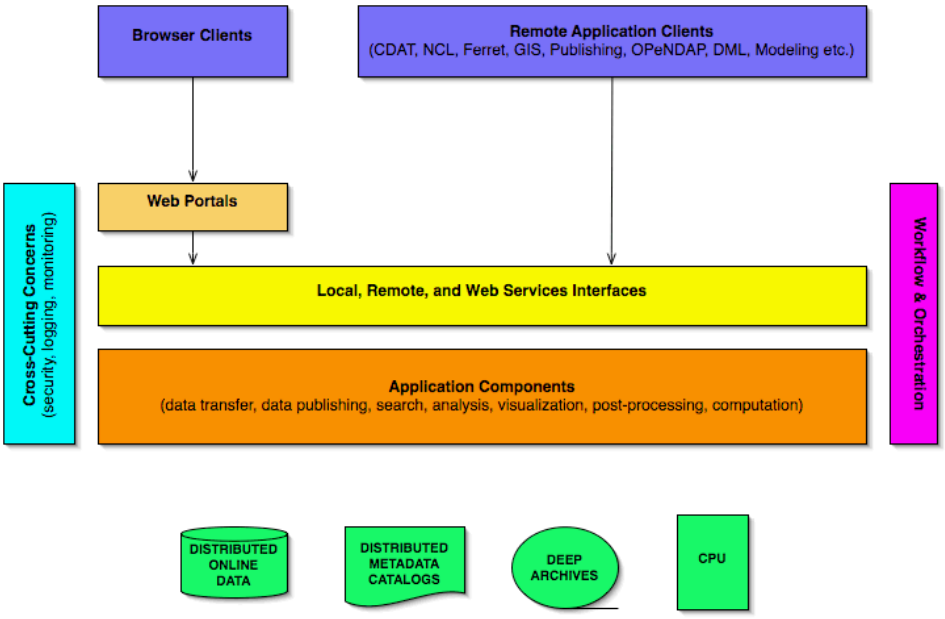

Figure 2: Future ESG-CET Architecture

The architecture must also be robust in the face of system and network failures at the participating sites. To address these concerns, we redesigned the federated ESG-CET architecture (see Figures 2 and 3) to provide interoperability and enhanced functionality to users [1]. The new architecture is based on three tiers of data services, as follows.
Tier 1 is a set of global services (partially exposed via a Global ESG Portal) that provide shared functionality across the overall ESG-CET federation. These services will include 1) user registration and management, 2) common security services for access control to the federated resources [13], 3) metadata services for describing and searching the massive data holdings, 4) common notification and registry services, and 5) global monitoring services. Through the global user management and security services, users will be able to access the whole system with a single set of credentials, and data owners will be able to control access as appropriate. The global metadata services will allow users to find data of interest throughout the federation, independently of the specific site at which a search is begun.

Various techniques will be used to provide levels of fault tolerance according to the requirements of the specific service.

Tier 2 comprises a limited number of ESG Data Gateways that act as data-request brokers for specific communities, possibly also providing specific enhanced functionality and/or documentation. Services deployed on a Gateway will include the user interface for searching and browsing metadata, for requesting data (including analysis and visualization) products, and for orchestrating complex workflows. We expect the software stack to be deployed at the Gateway level will require considerable expertise to install and maintain, and thus we envision these Gateways being operated directly by ESG-CET engineering staff. Initially, three ESG Gateways will be set up: one at LLNL to serve the IPCC AR4 and AR5 needs, one at ORNL to serve the Computational Climate End Station (CCES) project, and one at NCAR to serve the CCSM and PCM model communities (and possibly also others).

Tier 3 is the actual data holdings and the back-end data services that provide access to that data. These components will reside on numerous federated ESG nodes, each of which will typically host a suite of data and metadata services that are necessary to publish the data onto ESG, and to execute data-product requests formulated through an ESG Gateway. Because researchers and engineers at local institutions with varying levels of expertise will operate ESG Nodes, the mandatory software stack to be deployed at an ESG Node will be kept to a minimum, and supplied with detailed and exhaustive documentation. Sites may, however, deploy more extensive data processing capabilities if they wish, and the architecture allows another ESG site to retrieve and process data from a node lacking the desired capability.

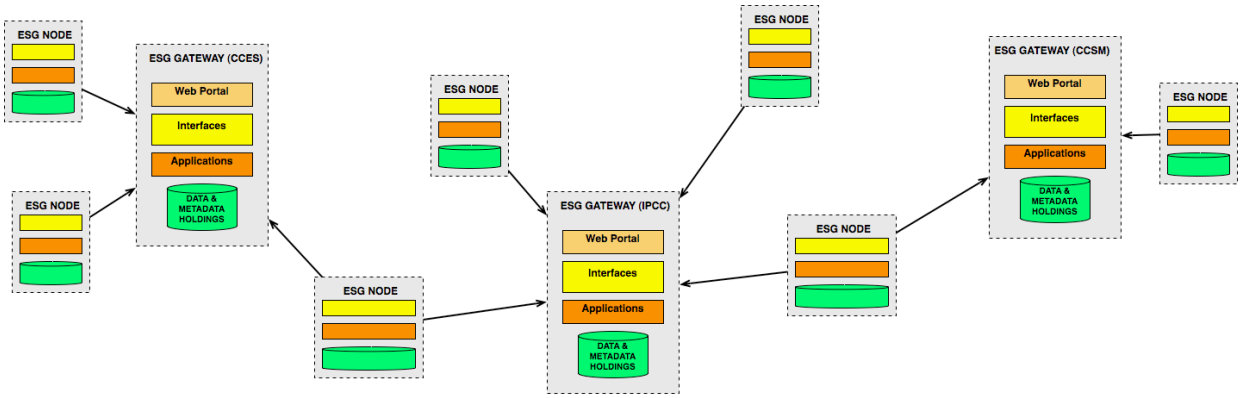

Figure 3: The ESG-CET Federated System

The proposed framework will support both individual ESG Nodes, configured to offer a core set of ESG functionality for local resources (data, metadata and computational resources), as well as a number of collaborating, thematic ESG Gateways (IPCC AR4 and AR5, CCES, CCSM, and others) that federate multiple ESG nodes and provide access to distributed resources.

We anticipate rolling out first prototypes and then robust production versions of these functions in sequence over the first four years of the project. Once a function is deployed, a combination of monitoring information and user feedback will be used to guide subsequent improvements to performance, robustness, and usability.

5. Conclusion

The Earth System Grid was initially developed and deployed under the SciDAC-1 program to meet the needs of the DOE's Climate Change Prediction Program to disseminate data to its research community. It also supported the international IPCC AR4 activity. Under SciDAC-2, the Earth System Grid Center for Enabling Technologies (ESG-CET) is carrying this work forward to meet the formidable challenges of the next phase of climate change research. Development efforts are underway to enable the deployment of a distributed infrastructure that is global in scope and that

supports petascale data volumes, more complex models and data, semantically based user interfaces, and a broader range of services, along with a variety of analysis and visualization applications.

6. Acknowledgements

This work is supported through the U.S. Department of Energy Office of Science, Offices of Advanced Scientific Computing Research and Biological and Environmental Research, through the SciDAC program. Argonne National Laboratory is managed by Argonne Chicago LLC under Contract DE-AC02-06CH11357. Information Sciences Institute is a research institute of the Viterbi School of Engineering at the University of Southern California. Lawrence Berkeley National Laboratory is managed by the University of California for the U.S. Department of Energy under contract No. DE-AC02-05CH11231. Lawrence Livermore National Laboratory is managed by the University of California for the U.S. Department of Energy under contract No. W-7405-Eng-48. Los Alamos National Security is managed by LLC (LANS) for the U.S. Department of Energy under the contract No. DE-AC52-06NA25396. National Center for Atmospheric Research is managed by the University Corporation for Atmospheric Research under the sponsorship of the National Science Foundation. Oak Ridge National Laboratory is managed by UT-Battelle, LLC for the U.S. Dept. of Energy under contract DE-AC-05-00OR22725. Pacific Marine Environmental Laboratory is under the National Oceanic and Atmospheric Administration's line office of Ocean and Atmosphere Research, lies within the U.S. Department of Commerce.

The ESG-CET executive committee consists of David Bernholdt, ORNL; Ian Foster, ANL; Don Middleton, NCAR; and Dean Williams, LLNL.

References

1. Earth System Grid Federation Architecture, http://esg-pcmdi.llnl.gov/documents/architecture/ESG-CET_architecture_20070119.pdf/, 2007.
2. Intergovernmental Panel on Climate Change (IPCC) 4th Assessment Report, http://www.ipcc.ch/activity/ar.htm, 2007.
3. World Climate Research Program (WCRP) CMIP3 Subproject Publications, http://www-pcmdi.llnl.gov/ipcc/subproject_publications.php, 2007.
4. Allcock, B., Bresnahan, J., Kettimuthu, R., Link, M., Dumitrescu, C., Raicu, I. and Foster, I., The Globus Striped GridFTP Framework and Server. in *SC'2005*, (2005).
5. Bernholdt, D., Bharathi, S., Brown, D., Chanchio, K., Chen, M., Chervenak, A., Cinquini, L., Drach, B., Foster, I., Fox, P., Garcia, J., Kesselman, C., Markel, R., Middleton, D., Nefedova, V., Pouchard, L., Shoshani, A., Sim, A., Strand, G. and Williams, D. The Earth System Grid: Supporting the Next Generation of Climate Modeling Research. *Proceedings of the IEEE*, *93* (3). 485-495. 2005.
6. Chervenak, A., Deelman, E., Foster, I., Guy, L., Hoschek, W., Iamnitchi, A., Kesselman, C., Kunst, P., Ripenu, M., Schwartzkopf, B., Stockinger, H., Stockinger, K. and Tierney, B., Giggle: A Framework for Constructing Scalable Replica Location Services. in *SC'02: High Performance Networking and Computing*, (2002).
7. Chervenak, A., Schopf, J.M., Pearlman, L., Su, M.-H., Bharathi, S., Cinquini, L., D'Arcy, M., Miller, N. and Bernholdt, D. Monitoring the Earth System Grid with MDS4 *2nd IEEE Intl. Conference on e-Science and Grid Computing (e-Science 2006)*, Amsterdam, Netherlands, 2006.
8. Collins, W.D., Bitz, C.M., Blackmon, M.L., Bonan, G.B., Bretherton, C.S., Carton, J.A., Chang, P., Doney, S.C., Hack, J.J., Henderson, T.B., Kiehl, J.T., Large, W.G., McKenna, D.S., Santer, B.D. and Smith, R.D. The Community Climate System Model Version 3 (CCSM3). *J. Climate*, *19* (11). 2122-2143. 2006.
9. Foster, I., Alpert, E., Chervenak, A., Drach, B., Kesselman, C., Nefedova, V., Middleton, D., Shoshani, A., Sims, A. and Williams, D., The Earth System Grid: Turning Climate Datasets

Into Community Resources. in *82nd Annual American Meteorological Society Meeting*, (Orlando, FL., 2002).

10. Foster, I., Nefefobva, V., Ahsant, M., Ananthakrishnan, R., Liming, L., Madduri, R., Mulmo, O., Pearlman, L. and Siebenlist, F. Streamlining Grid Operations: Definition and Deployment of a Portal-based User Registration Service. *J. Grid Computing*, *4* (2). 135-144. 2006.
11. Fox, P., Garcia, J. and West, P. OPeNDAP for the Earth System Grid, in preparation. *Data Science Journal*. 2007.
12. Shoshani, A., Sim, A. and Gu, J. Storage Resource Managers: Essential Components for the Grid. in Nabrzyski, J., Schopf, J.M. and Weglarz, J. eds. *Grid Resource Management: State of the Art and Future Trends*, Kluwer Academic Publishers, 2003.
13. Welch, V., Siebenlist, F., Foster, I., Bresnahan, J., Czajkowski, K., Gawor, J., Kesselman, C., Meder, S., Pearlman, L. and Tuecke, S., Security for Grid Services. in *12th IEEE International Symposium on High Performance Distributed Computing*, (2003).