SANDIA REPORT

SAND2004-1024
Unlimited Release
Printed March 2004

ASCI Grid Services Summary Report

Kathie L. Hiebert-Dodd

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation,
a Lockheed Martin Company, for the United States Department of Energy's
National Nuclear Security Administration under Contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: 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, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent 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, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865)576-8401
Facsimile: (865)576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: http://www.doe.gov/bridge

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone: (800)553-6847
Facsimile: (703)605-6900
E-Mail: orders@ntis.fedworld.gov
Online order: http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online

SAND2004-1024
Unlimited Release
Printed March 2004

ASCI Grid Services

Summary Report

Kathie L. Hiebert-Dodd
Distributing Computing Department
Sandia National Laboratories
PO Box 5800
Albuquerque, NM 87185-1137

Abstract

The ASCI Grid Services (initially called Distributed Resource Management) project was started under DisCom² when distant and distributed computing was identified as a technology critical to the success of the ASCI Program. The goals of the Grid Services project has and continues to be to provide easy, consistent access to all the ASCI hardware and software resources across the nuclear weapons complex using computational grid technologies, increase the usability of ASCI hardware and software resources by providing interfaces for resource monitoring, job submission, job monitoring, and job control, and enable the effective use of high-end computing capability through complex-wide resource scheduling and brokering. In order to increase acceptance of the new technology, the goal included providing these services in both the unclassified as well as the classified user's environment.

This paper summarizes the many accomplishments and lessons learned over approximately five years of the ASCI Grid Services Project. It also provides suggestions on how to renew/restart the effort for grid services capability when the situation is right for that need.

Acknowledgements

Over the years, many people have worked on the Grid Services project, and I want to thank all of them for their dedication and hard work. I especially want to thank the current team: Esther Baldonado, Daniel Fellig, Brian Jones, Fran Kanipe, Lois Lauer, Pat Moore, Mark Smith, and Shirley Starks, who helped make the decision to downsize. It was not an easy decision, but I believe it was the right decision.

I appreciate their willingness to develop alternative solutions for our customers and to leave the project in a position that, when a Tri-lab and/or SNL Grid is attempted again, the new developers will not have to start from scratch. I also want to thank the current team, as well as Steven Humphreys and Mike Sjulin, for their comments and suggestions about this report, and Kelly Gomez for her help in preparing the report.

I want to thank Mike Vahle, Mike Sjulin and John Noe for their support.

Contents

Figures

This page left blank intentionally

ASCI Grid Services
Summary Report

I. Introduction

This paper summarizes the many accomplishments and lessons learned over approximately five years of the ASCI Grid Services Project. It also provides suggestions on how to renew/restart the effort for grid services capability when the situation is right for that need.

The ASCI Grid Services (initially called Distributed Resource Management) project was started under DisCom² when distant and distributed computing was identified as a technology critical to the success of the ASCI Program[1]. The goals of the Grid Services project has and continues to be to provide easy, consistent access to all the ASCI hardware and software resources across the nuclear weapons complex using computational grid technologies, increase the usability of ASCI hardware and software resources by providing interfaces for resource monitoring, job submission, job monitoring, and job control, and enable the effective use of high-end computing capability through complex-wide resource scheduling and brokering. In order to increase acceptance of the new technology, the goal included providing these services in both the unclassified as well as the classified user's environment.

The Grid Services project has been a very successful project in terms of developing and deploying production quality capability. It successfully deployed the initial capability into the ASCI classified network, providing access to ASCI Red, ASCI Blue Mountain, and most importantly ASCI White. As part of DisCom², the Grid Services project met the ASCI Milepost with "flying colors". The Milepost External Review Panel complimented the team for the work that they had accomplished in providing a Kerberos version of Globus for the larger Grid community. The team has also deployed a Sandia-only unclassified grid that provides access to Janus, CPlant™, CPlant™ CA, and several other resources and planned to add the institutional computers.

The challenge has been, and continues to be, attracting and maintaining a significant number of users.

II. Grid Services Software Architecture and Deployed Configuration

The Grid Services software architecture provides a layered set of services. (See Figure 1) This approach was taken because it addressed the following ASCI/DisCom² related challenges:

---

[1] ASCI Technology Prospectus, A publication of the ASCI Program, National Nuclear Security Administration, Department of Defense, Defense Programs, DOE/DP/ASC-ATP-001, SAND2001-1765P, Sandia National Laboratories, Albuquerque, NM, July 2001, p. v.

- the solution must meet or exceed NNSA's security policies;
- any services provided must work side-by-side with the present set of remote computing services; and
- the solution must allow local sites to maintain control of their resources.

After a survey of several alternatives that were available at the time, early 1999, the Globus Toolkit[2] was chosen as the layer because it could provide the following basic grid services:

- a job submission protocol independent of the local resource manager;
- a standard output and standard error redirection mechanism;
- job control and status reporting; and
- health and status monitoring.

The Grid Services team modified Globus in two critical areas:

a) to use NNSA approved Kerberos V5 for authentication between Globus components; and
b) to interface with ASCI specific local resource managers.

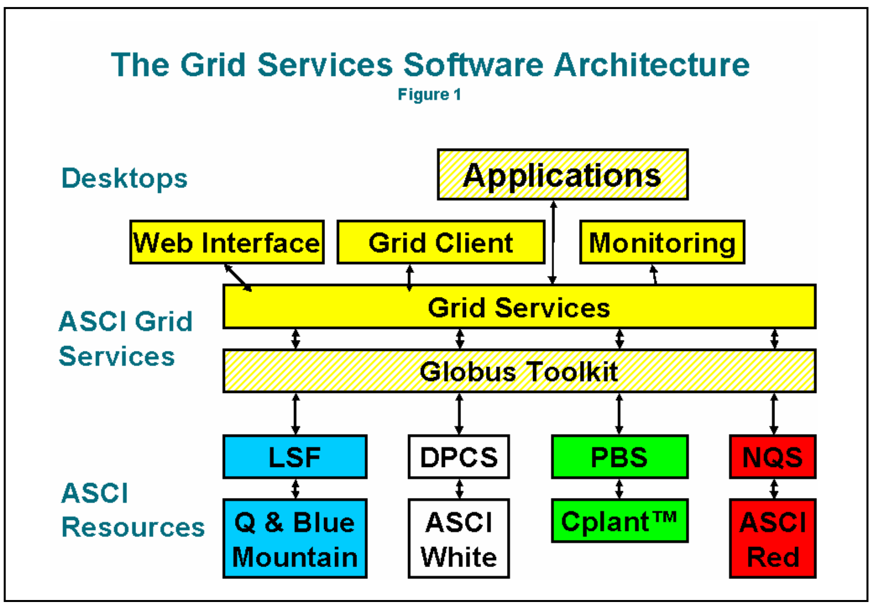

Figure 1: The Grid Services Software Architecture.

[2] Globus Toolkit, Argonne National Laboratories, University of Southern California Information Sciences Institute, and the Aerospace Corporation. URL://www-unix.globus.org/toolkit/

Figure 2 shows the configuration of the ASCI classified grid that was deployed in March, 2001, to meet the ASCI Level 1 Milestone. This configuration has not changed. Additional computer resources, most notably is LANL's Q machine, have been added. Sandia and LANL chose to have independent LDAP resource databases to provide resource information and better fault tolerance and redundancy. LLNL chose to use Sandia's LDAP.

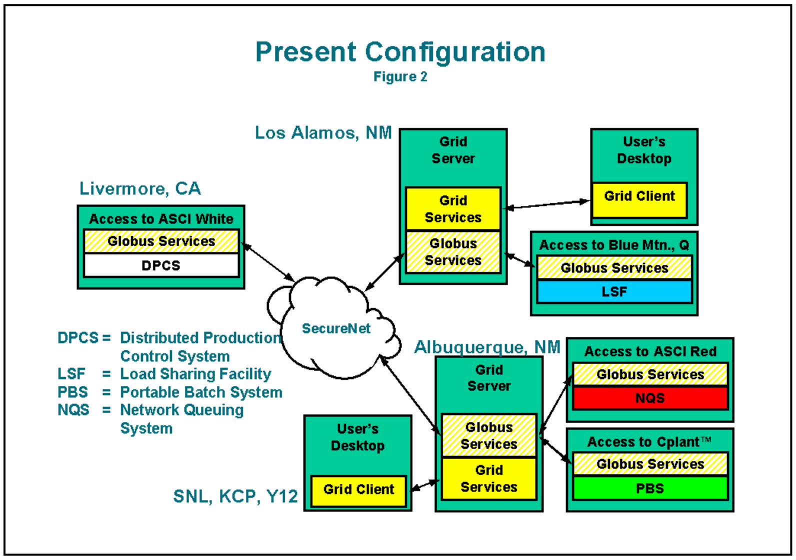

Figure 2: Present Configuration.

The Grid Services team also deployed a Sandia-Only Unclassified grid that included the following resources: Janus (unclassified ASCI Red), Ross (unclassified CPlant), Tesla (corporate SGI), phobos (viz cluster), and several machines on the Science and Engineering LAN (electron, maverick, rogue, and the SGI cluster). Because Janus is on the HPC LAN, which is a part of the SON, and the remainder of the resources is on the SRN, two Grid servers and LDAPs were deployed for this Grid.

III. Deployed Capabilities

By March 2001, for the DisCom^2 Milepost, the following capabilities were deployed to the ASCI Tri-lab classified high performance computing environment:

1. Desktop tools that allow users to submit, monitor, and control jobs on both local and remote resources without a single remote login.
2. Workflow management that supports serial job dependencies, file staging, and standard output and standard error redirection.
3. Brokering and resource discovery that locates hardware and software resources on behalf of the user and selects the "best" available resource.
4. Job submission protocols independent of the resource's local scheduler.
5. Web-based monitoring that allows users to access resource and job status from all three nuclear weapons laboratories from a single web page.

The goal was to enhance this initial set of grid services, which were usable, into a set of services that users would find indispensable.

Figure 3 illustrates the Production Wizard, the initial desktop tool for submitting, monitoring and controlling workflow. It required the manual creation of a Gale[3] XML workflow script using a standard text editor.

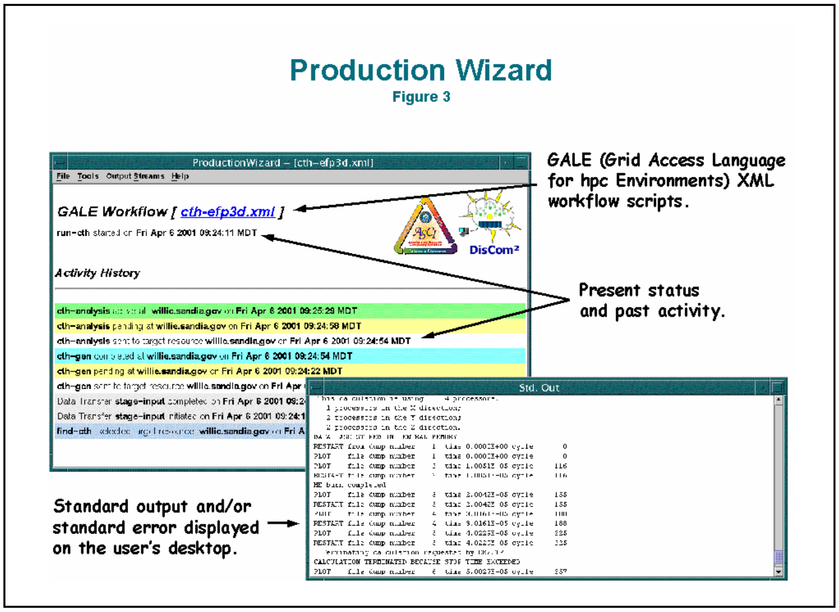

Figure 3: Production Wizard.

Figure 4 illustrates the early version of the web-based monitoring.

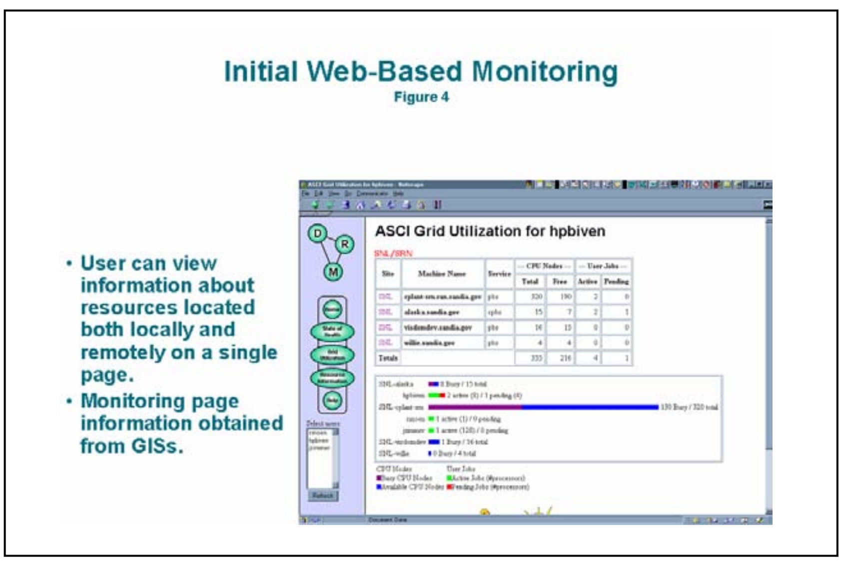

Figure 4: Initial Web-Based Monitoring.

---

3 "GALE: Grid Access Language for hpc Environments", Hugh Bivens and Judy Beiriger, 4/16/2001; https://wfsprod01.sandia.gov/groups/srn-uscitizens/@snl-employees/documents/document/wfs024016.pdf

Early ambitious plans[4] envisioned the following additional capabilities:

- Graphical workflow editor (Production Wizard 2 – see Figure 5, deployed but decommissioned beginning of FY03)
- Enhanced fault detection and reporting (done, but could be improved)
- More sophisticated brokering algorithms (This research was not encouraged by the ASCI community. Only recently has there been interest.)
- Parallel WAN data transfer support (done -- Zephyr/Cyclone)
- User credential refresh mechanisms (done)
- Administrative support tools to automatically add grid users (done; also use of Webcars to request Grid account)
- Support parallel workflow (not attempted)
- Develop automated and intelligent job restart mechanisms (only preliminary prototypes were developed).
- Grid-enable non-compute resources such as: archival storage, data services, and visualization (used to access HPSS, sample workflows to use Ensight and Server of Servers)
- Coordinate the use of multiple resources (some research and development, never deployed)

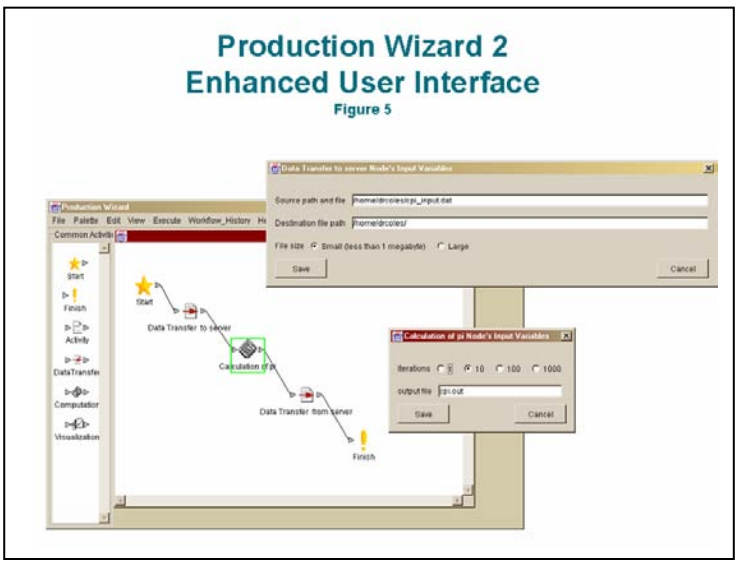

Figure 5: Production Wizard 2 – Enhanced User Interface.

The following capabilities, available in September 2003, were deployed to and supported on corporate resources:

1. Web-based monitoring that allows users to access resource and job status from all three nuclear weapons laboratories from a single web page. (See Figures 6, 7, and 8. **Note**: only sample data is used in the figures.)

[4] ASCI DisCom[2] Milepost Presentation, Steven Humphreys, April 12, 2001; https://wfsprod01.sandia.gov/groups/srn-uscitizens/documents/document/wfs098481.pdf

2. On SRN, click on "Grid Services" in the "Pick an Engineering Resource" pull-down menu.
3. On SCN, click on "Grid Services" under Engineering Tools.

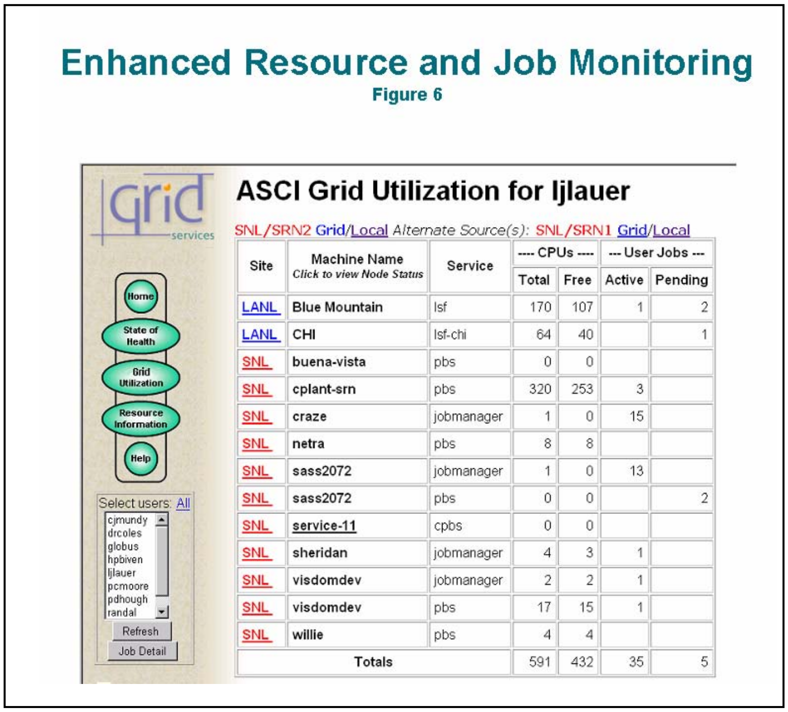

Figure 6: Enhanced Resource and Job Monitoring.

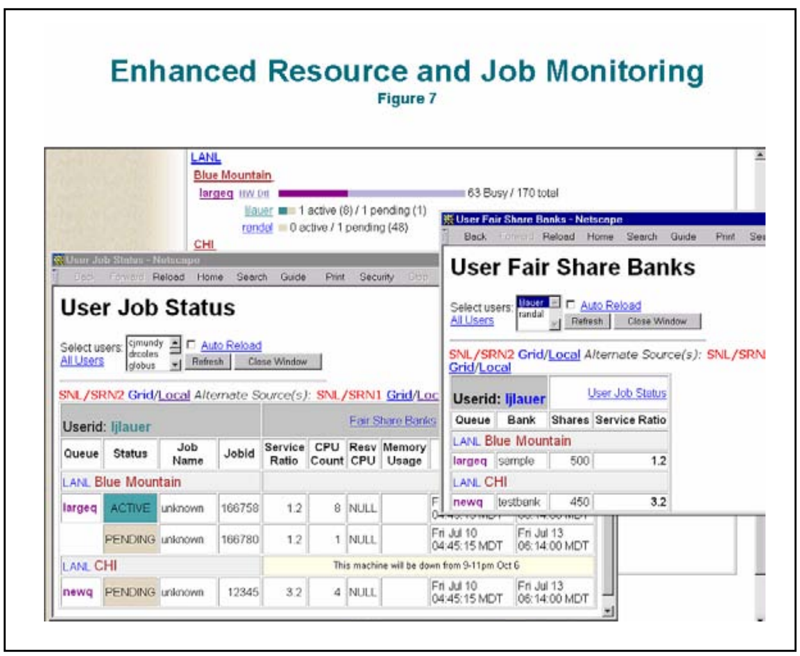

Figure 7: Enhanced Resource and Job Monitoring.

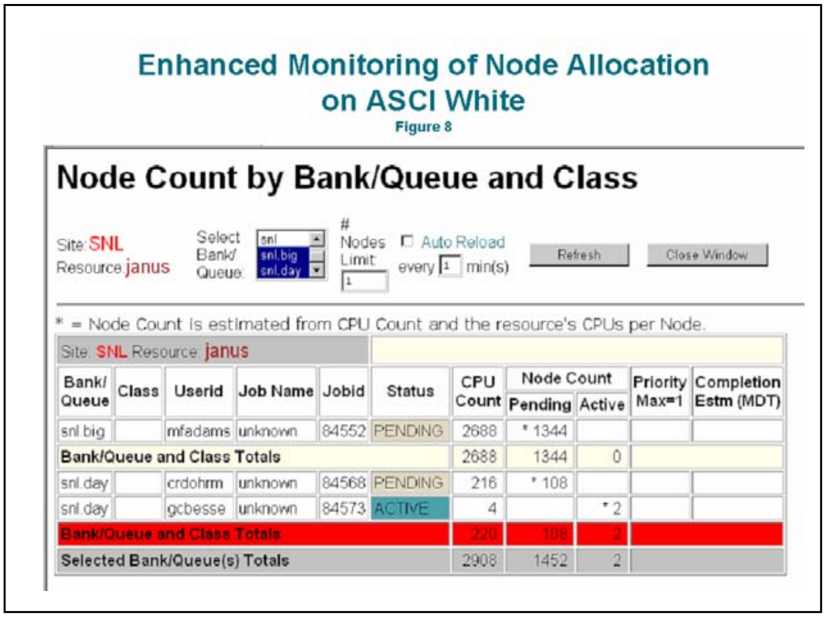

Figure 8: Enhanced Monitoring of Node Allocation on ASCI White.

4. Desktop and Web-based tools that allow users to submit, monitor, and control jobs on both local and remote resources without a single remote login. (See Figure 9)

Figure 9: Web-based Job Submission.

5. Persistent workflow management that can resume running workflows when systems fail. Existing features are still supported, such as serial job dependencies, file staging, and standard output and standard error redirection. Web-based monitoring of workflow status. (See Figures 10, 11, and 12)
6. Zephyr/Cyclone – a fault-tolerant file transfer capability.

7. Brokering and resource discovery that locates hardware and software resources on behalf of the user. (No change from original deployment.)
8. Webcars request for “Grid Account”.
9. DRM-Help e-mail on both SRN & SCN.
10. Grid Services are automatically monitored and alarms sent to Sitescope.

Figure 10: Web-based Job Monitoring.

Figure 11: Web-based Job Monitoring.

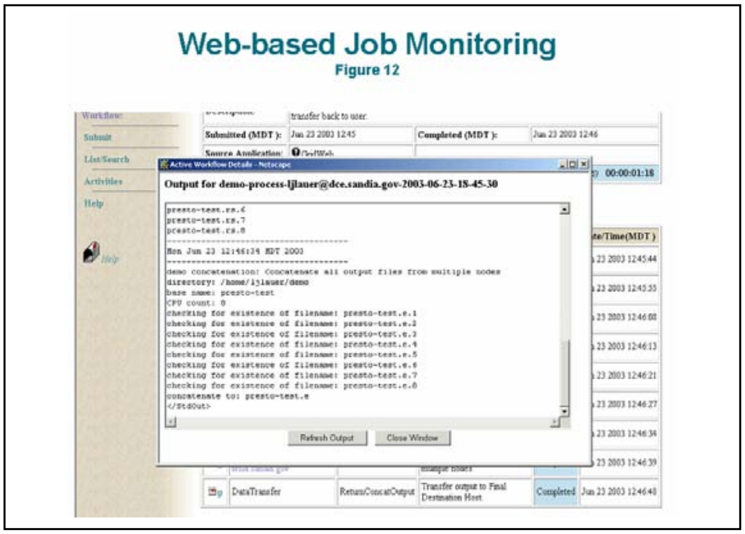

Figure 12: Web-based Job Monitoring.

The main differences between March 2001 and September 2003 is a tighter integration with corporate computing, features to address fault-tolerance and robustness issues such as persistent workflow, additional information about resources, e.g. node count, information about the status of the workflow, Zephyr/Cyclone, automated monitoring, and better user support.

IV. Other Accomplishments

A. Grid Community Participation

Until October 2002, Grid Services team members took an active role in first the Grid Forum and then the Global Grid Forum. In early 2001, the Grid Forum merged with eGrid (European Grid Forum) and the Asia-Pacific Grid Forum to create the Global Grid Forum. They attended and participated in all the Grid Forum meetings, 1 - 5 and the first 5 Global Grid Forums. They took lead roles in the following Working/Research Groups: Grid Computing Environments, Grid User Services, Dictionary, Grid Information Services, Grid Security, Scheduling and Resource Management and the Steering committee that plans the GGF meetings. They contributed to the development of standards in Grid Security, User Services, and Grid Dictionary.

B. Open Source Distribution

With support from the ASCI NNSA Headquarters[5], the Grid Services software, DRM version 1.3[6] was copyrighted and then distributed as open source under the lesser GNU license. Prior

[5] "Release of Open Source Software", William Reed, Oct. 2001; https://wfsprod01.sandia.gov/groups/srn-uscitizens/documents/document/wfs152210.pdf

to the open source distribution, the code was distributed to Argonne, NASA, and DoD via a Government Use Only license.

C. Platform Contract

As of September 2002, the Grid services software was based on old versions of Globus (v1.3), and LDAP software, and was facing a major funding cut for FY03. While the Globus Toolkit is open source, Platform was advertising that they could provide productionization of the code, custom changes, and maintenance support. It was decided to contract with Platform to provide the necessary security changes to Globus Toolkit Version 2 based on Sandia's original work. The "kerberized" Globus Toolkit Version 2 was to be evaluated, and if warranted, the ASCI grids would be upgrade to the new version. Because of delays in delivery from Platform, several bugs discovered in the delivered product, and subsequent decisions leading to a downsizing of the effort, the Platform software was never deployed.

Platform has given their changes that resulted from this contract back to the Grid community.

D. Grid-Enabled Applications

The Grid Services team has worked with the following ASCI S&CS and CSRF tool development teams to "grid-enable" their tools (provide access to the grid computational resources): Data Services, Simba, SimTracker, Entero (ITS & Xyce), Heidi/Chilespice, 9200's Design Simulators, and Dakota. The Data Services tools, SimTracker, and Heidi/Chilispice use a web-based interface. The Entero partnership was the most successful. Grid-enable Entero was used to support the FY03 ASCI Capabilities Level 1 Milestone. Prototypes were developed for the other projects, but for numerous reasons, many times the application wasn't ready to deploy, production grid-enabled versions were not deployed.

V. We have built it, but they are not coming!

Fairly early in the deployment of the Grid Services, the monitoring capability was used and continues to see moderate use by ASCI users, ~1400 hits over the last quarter of FY03. Late FY02, Grid Services developed Cyclone (now Zephyr), a robust file transfer capability that was built on top of the existing Grid Services. Cyclone automatically transfers a directory of large files from White to Sandia's HPSS system, performs checksums to verify the accuracy of the transfers, monitors the transfer, and for files that didn't transfer successfully, retries several times. Cyclone transferred over 175,000 files and 9.5 terabytes of data from White. The potential of this capability brought enough show of support from users, that SNL was able to persuade LLNL to keep White on the Grid.

[6] ESTSC License Form, Jan. 2002; https://wfsprod01.sandia.gov/groups/srn-uscitizens/documents/document/wfs152209.pdf

The Grid Services team has performed numerous requirements gathering efforts.[7,8] A User Survey was conducted in summer 2001[9].The team has worked very hard with individual users and tool developers, but we still don't have users routinely using the Grid to submit their jobs.

VI. Why is it so hard?

The ASCI Tri-lab High Performance Computing environment is a very complex environment. The ASCI Grid is comprised of a few highly unique and temperamental supercomputing platforms, rather than a large number of relatively stable workstations and small clusters. There are no "standards" and very little coordination and cooperation among the three labs with respect to even basic services, e.g. the version of SSH, other software upgrades, configuration of the machines, security mechanisms, queuing systems, I/O, usage policy…. The Globus Toolkit was able to abstract some of the heterogeneity of the environment, but not enough for the Tri-lab environment. The goal of "write once; run anywhere" was never completely achieved.

This point is important enough to bear repeating: building grid systems on top of heterogeneous and un-standardized platform environments is very difficult. Differences in versions of core software such as SSH, Kerberos, and PFTP, as well as system configurations (location of home directories, cross-mounting policies, file storage access restrictions, etc), pose too much complexity to reasonably handle in a grid abstraction layer. By itself, Globus does not provide enough uniformity among systems to provide a solid standard base to build upon. A "Grid COE" for scientific computing standard across the Tri-lab could have provided this necessary base. For example, the ATLAS[10] grid project requires all systems to install a package called VDT (Virtual Data Toolkit)[11], which includes a standard set of tools that their grid software can depend on.

ASCI size problems -- basically the size and number of the files and the length of the calculations cause problems. The team has found bugs in basic services such as PFTP and checksum.

The Grid Services team developed the concept of a workflow, a linear sequence of job steps. This current capability doesn't fit the needs of our potential users – basically it is not at the appropriate level of granularity and it doesn't provide enough control. Additionally, the workflow is very different from the current way staff interfaces with the platforms, i.e. command line. It appears that it was too big a step.

[7]" Distance and Distributed Computing and Communication FY99 Service Model", Judy I. Beiriger; Ann L. Hodges; Wilbur R. Johnson; Martha J. Ernest; Robert D. Pollock, SAND99-0683, 10/30/99, https://wfsprod01.sandia.gov/groups/srn-uscitizens/documents/document/wfs026212.pdf

[8]"FY03 Elicited Requirements Analysis, May, 2003; https://wfsprod01.sandia.gov/groups/srn-uscitizens/documents/document/wfs075267.pdf

[9] Grid Services Survey, https://wfsprod01.sandia.gov/groups/srn-uscitizens/documents/document/wfs047969.pdf

[10] ATLAS; http://www.usatlas.bnl.gov/computing/grid/

[11] Virtual Data Toolkit; http://www.lsc-group.phys.uwm.edu/vdt/

Advertising the Grid Services has been left to the team because there has been no central customer service to help users and suggest using the Grid.[12] For individual platforms and some specific codes, there are efforts to help users. But, to date, these efforts have not included the Grid Services as a way of submitting jobs.

The intended "user" has been a moving target. Most recently, the DART (Design thru Analysis) Team has been trying to represent the user community. Defining what would make users more productive is difficult for both users and developers to determine. As we have worked with individual users, they have expected us to be experts in everything – not just accessing the platforms, but every aspect of the platform as well as experts in the particular capability code. The Grid developers have to reverse-engineer scripts to understand how a user is running a code to understand how to develop a workflow. Recently, a user wanted a team member to port their code to a new platform and get it running for them. Additionally, many users are willing to spend the time to learn how to run on a new platform. These users are also willing to share their "point solution" with others, thus propagating an inefficient process that generates more undocumented, unsupported, unadvertised code, and removes yet another opportunity for the benefits of Grid Services to be realized.

What is the appropriate user interface? The initial Production Wizard was dismissed because "users" would have to develop a Gale XML script. Many believed that a GUI that would create the script for the user had to be the only interface. This development effort was costly and never used because of its lack of flexibility. It became apparent that there was no "one size fits all" interface solution; some analysts are very familiar with the Unix command-line environment and are unwilling to sacrifice control for a GUI, while others would actually prefer a graphical interface. Finding a happy medium between the two extremes is a difficult challenge for new software products in the scientific world. With a customer base that spans a large range of interest and willingness to learn to run on the new machines, interface solutions must support both the "old-school" and newer computing paradigms, be easy-to-use, and provide a lot of customization, tuning, and integration with existing solutions. In addition, users tend to become comfortable using a particular set of tools, and in order to switch to something new the perceived benefits must outweigh the perceived costs. The uncertainty of dealing with new grid technologies, which might abstract away important details, is a large perceived cost for users. An extra layer of software also adds more failure modes and presents a new set of debugging challenges when something goes wrong.

Working in the Tri-lab environment has been difficult. LLNL never believed in the Grid vision Initially, they didn't need to believe because White was not a remote resource. LANL's support was concentrated in a single advocate, who's funding to support the Grid was eliminated at the beginning of FY03. In FY03, ASCI Headquarters moved the project to On-Going without any funding. The project was able to continue because of the support of Mike Vahle, Mike Sjulin and John Noe.

[12] In preparation for ASCI Red Storm, Sandia is developing a High Performance Computing (HPC) User Help Desk. This is a centralized service to help: 1.) Sandia users to use the big platforms and 2.) tri-lab users to use ASCI Red Storm.

VII. ASCI Grid Status January 2004

The Grid Service layer (see Figure 1) will be decommissioned on Feb.1, 2004. It was determined that the Globus layer should be kept for at least the remainder of FY04. Because of the users, primarily the monitoring and a few tool developers that are dependent on the Grid Services, the team developed the following alternative approaches:

- **Monitoring Web site:** only removed unsupported information, basically remains unchanged, depends on Globus, GIS services, and the LDAP database.
- **Design Simulators**: a set of "standardized" shell scripts for specific platforms, Tesla and White that does not use the Grid Services or Globus to submit a job.[13]
- **Heidi/Chilispice**: developed a servlet and a set of "standardized" shell scripts for specific platforms, Ross, Tesla, Electron and Maverick that enable remote job submission based on SSH and SCP and that does not use the Grid Services or Globus.[14]
- **Data Services:**
    - Zephyr/Cyclone: rewritten to use only the existing SSH and PFTP tools; does not depend on Grid Services or Globus.[15]
    - Job submission: developed a simplified job submission tool that depends only on Globus and not Grid Services.[16]

VIII. Recommendations

The number one recommendation is to build a coalition of platform, queue, and code experts that works closely with the users to develop a consistent, easy access to job submission solution. It takes the knowledge and cooperation of all three groups to understand the requirements and develop an **accepted** solution.

Therefore, the second recommendation is: in order to develop an **accepted** user interface, one needs to start simple, close to what the user is doing today, e.g. command-line. The interface must provide control and more importantly detailed information about status and what went wrong. Instead of replacing the existing use paradigms, grid tools must first integrate with them and augment them in easily identifiable ways. For example, instead of providing a web page GUI for file transfers, perhaps provide simple command-line tools that look and feel like the ones in use (SCP, PFTP). To support the workflow concept, instead of introducing a new XML-based language for flow description, perhaps provide grid command-line utilities which can be used in scripting languages the users already know such as perl, shell, or python. A key point is to keep the gap between new grid software and

---

[13] Design Simulators - Common Tesla Script, December, 2003; https://wfsprod01.sandia.gov/groups/srn-uscitizens/documents/document/wfs099813.pdf

[14] Heidi – Job Submission to Remote Hosts; https://wfsprod01.sandia.gov/groups/srn-uscitizens/documents/document/wfs152582.pdf

[15] ZephyrFTG Design Document; https://wfsprod01.sandia.gov/groups/srn-uscitizens/documents/document/wfs097723.pdf

[16] Data Services – Alternate Grid Services Integration; https://wfsprod01.sandia.gov/groups/srn-uscitizens/documents/document/wfs152776.pdf

existing techniques small each time new tools and functionality are introduced, and work towards more advanced grid visions in baby steps instead of giant ones.

Continue to work with other tool developers, such as ASCI PSE and CSRF that are developing customized tools and GUIs. Provide simple, reliable and robust grid utilities that these developers can easily incorporate into their software.

Promote a COE for scientific computing to standardize on core software such as SSH, Kerberos, and PFTP, as well as system and desktop configurations.

Finally, given human nature and the resistance to change, there will be the need to use “carrots” or “sticks”, e.g. Advanced Deployment funds to encourage users to use the Grid technology or a policy that would force Grid technology to be the only access to certain highly desirable machines, to help users adopt the new technology.

Appendix

People involved in the Grid Services Project:

Grid Services Development Team

Esther Baldonado	Judy Beiriger
Hugh Bivens	Dwight Coles
Daniel Fellig	Steven Humphreys
Wilbur Johnson	Brian Jones
Fran Kanipe	Ellen Lemen
Pat Moore	Ron Rhea
Mark Smith	Shirley Starks
Young Seiden	Mary Roehrig
Ruthe Vandewart	Marilyn Warrant

Nuclear Weapons Complex Team

William Barber, LANL
Randal Rheinheimer, LANL
Moe Jette, LLNL
Leigh Shoulta, KCP
Rhonda MacIntyre, Y12

PEARLS
Ann Hodges
Bryan Spicer

Management Team
Larry Ellis
Kathie Hiebert-Dodd

ASCI/DisCom2 Management Team
Art Hale
Michael Sjulin

ASCI/On Going Computing Management Team
John Noe

Distribution:

1	MS-0139	Mike Vahle, 9900
1	MS-0139	Robert K. Thomas, 9904
1	MS-0316	Joseph Castro, 9233
1	MS-0318	Paul Yarrington, 9230
1	MS-0321	Jim Ang, 9220
1	MS-0321	R.W. Leland, 9220
1	MS-0525	Steven Wix, 1734
1	MS-0525	Ron Sikorski, 1734
1	MS-0576	Hugh Bivens, 5536
1	MS-0797	Martha Ernest, 9333
1	MS-0801	Art Hale, 9300
1	MS-0801	Michael Sjulin, 9330
1	MS-0806	Leonard Stans, 9336
1	MS-0807	John Noe, 9328
1	MS-0820	Pat Chavez, 9232
1	MS-0822	Constantine Pavlakos, 9326
1	MS-0822	David White, 9227
1	MS-0822	Rena Haynes, 9227
1	MS-0841	Thomas Bickel, 9100
1	MS-0847	Ted Blacker, 9226
1	MS-0924	Dwight Coles, 9524
1	MS-1137	Mark Gonzales, 6224
1	MS-1137	Kathie Hiebert-Dodd, 6224
1	MS-1137	Esther Baldonado, 6224
1	MS-1137	Daniel Fellig, 6224
1	MS-1137	Steven Humphreys, 6224
1	MS-1137	Wilbur Johnson, 6224
1	MS-1137	Brian Jones, 6224
1	MS-1137	Fran Kanipe, 6224
1	MS-1137	Lois Lauer, 6224
1	MS-1137	Pat Moore, 6224
1	MS-1137	Mark Smith, 6224
1	MS-1137	Shirley Starks, 6224
1	MS-1140	Larry Ellis, 5502
1	MS-1179	Leonard Lorence, 15341
1	MS-1208	Judy Beiriger, 5536
1	MS-9012	Michael Hardwick, 8964
1	MS-9051	Young Seiden, 89452
1	MS-9018	Central Technical Files, 8945-1
2	MS-0899	Technical Library, 9616