

SciDAC DataGrid Middleware
A High-Performance Data Grid Toolkit:
Enabling Technology for Wide Area Data-Intensive Applications
Quarterly Report October 2001 thru December 2001

Accomplishments this Quarter:

- ***The Globus Toolkit Enters Public Beta Release:***

In November 2001 The Globus Toolkit V2.0TM entered public beta release. This followed a six month long, highly successful alpha release cycle. This release use the new packaging technology to allow for modular installation and updates of the toolkit, makes significant improvements to the Information Services (MDS), reliability enhancements to our Resource Management tools (GRAM), and the addition of new DataGrid tools (GridFTP for high speed, reliable transport and Replica Management for more efficient access to large datasets). For further information, please see:<http://www.globus.org/gt2/>

- ***Reliable File Transfer Demonstrated at SC 2001:***

During SC2001 in Denver, CO we demonstrated a prototype of our Reliable File Transfer (RFT) service. This was a Java based service which took a list of SourceURL, DestinationURL pairs and then moved them. Facilities were provided to submit, query, and cancel the jobs. The Globus GridFTP client library reliability and performance plug-ins were used to provide appropriate state data which was stored in a local database. In the event of a local failure, on restart the database was read and all in progress transfers were resumed and the outstanding queue restored. For information please see:<http://www.mcs.anl.gov/~madduri/RFT.html>

- ***GridFTP based Proxy Server Sinks 2.8 Gigabits/second during SC2001 Bandwidth Challenge:***

During the Bandwidth Challenge at SC 2001 in Denver, CO we simulated a large High Energy Physics site that was receiving data from multiple other sites. In order to support the high data rates necessary to achieve this, we prototyped a "proxy" server. The sites contacted a main server to initiate the transfer. This server then chose one of 16 nodes at its disposal to service the request. Data was being received from 27 hosts at 16 different sites at a peak rate of 2.8 Gbs, and a sustained rate of nearly 2 Gbs. For further information, please see: [SC2001 Bandwidth Challenge Entry](#)

- ***Provide secure remote I/O services to JAVA applications:***

As part of the new JAVA Universe of Condor, JAVA applications can use simple "POSIX like" I/O protocol to remotely read and write files via a an authenticated and secure communication channel provided on the executing machine. Initial tests demonstrated the effectiveness of the capability for JAVA jobs submitted to a heterogeneous collection of Grid resources via Condor-G. The new Universe provides the application a JAVA and Grid aware execution environment.

- ***Support for multi-route writes in Kangaroo:***

Added to the Kangaroo remote I/O system the ability to transfer a file written by an application to its destination concurrently via different routes. This allows Kangaroo to exploit all available

bandwidth and buffering capacity of a data Grid when moving the data from the production site to the storage site.

- ***Defining, Managing and Scheduling Data Placement (DaP) jobs:***

We developed a ClassAd based Job Description Language for Data Placement (DaP) jobs. The language allows jobs to define their storage requirements at the destination and the method to be used to place the data. We are in the process of integrating DaPs into our DAG framework and interfacing the DAGMan with a DaP scheduler. The scheduler is responsible for managing, executing and monitoring DaP jobs.

- ***Reliable and efficient management of ClassAd collections:***

Continued to enhance and harden our library for managing persistence collections of ClassAds. We added support for caching subsets of a collection in memory. A reliable and recoverable protocol for inserting, removing and updating ClassAds in a collection was defined and implemented. The protocol enables clients to remotely manipulate the content of a ClassAd collection.

- ***Develop a framework for Error and Exception handling in Grid I/O:***

Worked on developing a framework for managing the information and activity flow of error and exception handling in a layered, multi-vendor, distributed I/O system. Specifically, this effort address the following challenges:

1. Effective communication of the causes to a failed I/O operation.
2. Reliable logging and exchange of information about failed operation.
3. Methods to detect the causes of a failed operation.
4. Definition and execution of retry policies for failed operations
5. Exposing information on failed I/O operations to job schedulers and resource managers.

Plans for next quarter

- ***Continued Beta Testing for the Globus Toolkit:***

With the introduction of the new packaging system and shared libraries, we face some new challenges in our initial port to other platforms. We will continue to resolve issues related to operation on platforms beyond Linux and Solaris and of course resolve any bugs that are discovered.

- ***Continue Refinement of the Reliable File Transfer Service:***

We will continue our prototyping activities on RFT. This will include the addition of security, a web services front end, resolution of problems with large numbers of concurrent transfers of large files, and some initial stress testing.

- ***GridFTP improvements***

Work on GridFTP will include support of the Beta testing cycle and additions to the protocol and implementation. These additions are intended to resolve some of the limitations discovered in the protocol and provide improved functionality. We will also begin investigation and design discussions on a new data channel protocol which will allow bi-directional data flow, full duplex operation, and pipelining of commands.

- ***Continued Design Work and Prototypes of the new Replica Management System***
We have been working in conjunction with the EU DataGrid Project on the design of an improved Replica Management system. This is primarily to distribute the catalog to eliminate a single point of failure and to improve scalability. We will continue refining this design and expect to produce initial prototypes for evaluation of the current design.

Papers Published or in Progress

High-Performance Remote Access to Climate Simulation Data: A Challenge Problem for Data Grid Technologies. B. Allcock, I. Foster, V. Nefedova, A. Chervenak, E. Deelman, C. Kesselman, J. Leigh, A. Sim, A. Shoshani, B. Drach, D. Williams, SC 2001, November 2001.

Giggle: A Framework for Constructing Scalable Replication Services was presented in draft form to the EU DataGrid project for their input. Giggle is being developed as a joint effort with the EU DataGrid Project.

Local Name Resolution Catalog Specification was presented in draft form to the EU DataGrid project for their input. Giggle is being developed as a joint effort with the EU DataGrid Project.

Applied Techniques for High Bandwidth Data Transfers Across Wide Area Networks Jason Lee, Dan Gunter, Brian Tierney, Bill Allcock, Joe Bester, John Bresnahan, Steve Tuecke (*Proceedings of Computing in High Energy Physics (CHEP '01)*) September 2001

Globus Toolkit Support for Distributed Data-Intensive Science W. Allcock, A. Chervenak, I. Foster, L. Pearlman, V. Welch, M. Wilde (*Proceedings of Computing in High Energy Physics (CHEP '01)*) September 2001

Presentations Given

1. 28-31 Jan 2002: Globus Developers Tutorial, Argonne National Laboratory, Chicago, Illinois, USA
2. 21-24 Jan 2002: Globus Developers Tutorial, Edinburgh Scotland
3. 10 Jan 2002: Presentation of Latest Replica Catalog Design at PPDG Reliable Replication Meeting, Jefferson National Accelerator Laboratory, Newport News, Virginia, USA
4. 05 Dec 2001: "Replica Management for the Grid" presented at NASA Information Power Grid Workshop, Palo Alto, CA
5. 13 Nov 2001: A technical paper titled "High Performance Remote Access to Climate Simulation Data: A Challenge Problem for Data Grid Technologies" was presented at SC 2001, Denver, CO
6. 12-15 Nov 2001: Earth Science Grid Demo incorporating GridFTP and Replica Catalog presented at SC 2001, Denver, CO
7. 12-15 Nov 2001: Ligo Demos incorporating GridFTP and Replica Catalog presented at SC 2001, Denver, CO
8. 11 Nov 2001: DataGrid Tutorial presented at SC 2001, Denver, CO
9. 08 Oct 2001: "Giggle Framework Overview" presented at the Global Grid Forum meeting in Frascati, Italy.
10. 15 Oct 2001: Replica Catalog Breakout session at GriPhyN meeting, Marina Del Rey, CA

11. 15 Oct 2001: Virtual Data Catalog Breakout (uses GridFTP and Replica Management) session at GriPhyN meeting, Marina Del Rey, CA