

**SciDAC DataGrid Middleware**  
**A High-Performance Data Grid Toolkit:**  
**Enabling Technology for Wide Area Data-Intensive Applications**  
**Quarterly Report January 2004 thru March 2004**

**Accomplishments this Quarter:**

**Globus Toolkit V3.2 was released**

On March 31<sup>st</sup>, version 3.2 of the Globus Toolkit was released. This release was the first to contain the new Globus eXtensible Input/Output (XIO) system, an improved GridFTP server, and scalability improvements to the Reliable File Transfer (RFT) service. More details are below.

**eXtensible Input / Output (XIO) System**

The release of GT3.2 marks the first stable release of XIO. Every C based component is using XIO via a globus\_io wrapper. The beta and initial response from the release have been very positive. A few bugs have been found, but quickly fixed. With its ability to abstract away transport differences, provision for synchronous or asynchronous calls, and uniform error handling, it can greatly ease the transition to different transport protocols and/or data sources and sinks.

**wuftp based GridFTP**

The release of GT3.2 marks some long awaited improvements in our GridFTP implementation (it still remains backwards protocol compatible). This also marks the last feature enhancements that will be made to the wuftp based implementation. New features include:

- **MLSx commands:** MLST and MLSD are RFC defined commands for FTP that provides consistent structured directory listings that are programmatically parseable.
- **File Globbing:** globus-url-copy now supports moving multiple files with a single command. It can accept a directory or a glob (such as \*.dat) and will move all the files. All features of the protocol are functional and channel caching is used to improve performance.
- **Checksum Command:** Allows for end to end verification of the transfer by doing a checksum on the disk image on each end after a transfer. This can detect things that the TCP checksum might miss.
- **Switch for RFC1738 compliant URLs:** Most FTP servers, and our GridFTP server, do not follow the RFC when it comes to interpreting FTP URLs. For most situations, this is a good thing. There are some communities for whom this is a problem. We now provide a switch that will interpret the URLs according to the RFC
- **chmod Support:** The GridFTP server now adds chmod to its list of filesystem operations it can perform. Note that this is not standard and will have to be introduced to the GGF for standardization.

### **Status of new GridFTP server**

The new GridFTP server is nearly feature complete. We are planning an experimental release in late April or early May. This will be a complete protocol identical replacement for the wuftp based server. Installation and configuration are different, but any previous client that worked with a 2.4.x or later server will work with this server. This version will \*NOT\* include striping, but that will be our next feature addition after the experimental release.

### **Reliable File Transfer (RFT) Service**

With the release of GT3.2 RFT can handle approx 50 times as many files as the previous version (now approx 20,000 files). The downside to this is that there is currently no way to get the full request status. This required the addition of new port types and it was decided to not do this until the next GT release. Though it did not make the release deadline, we believe we have overcome a bug that was limiting scalability and can now handle over a million files, though we have not been able to verify this yet.

### **Replica Location Service Development**

We have continued developing and testing the Replica Location Service for the Globus Toolkit Version 3.2. The development included the release of an RLS version that supports a hierarchical index of RLS servers. The latest version of the RLS code also includes some important bug fixes, including a memory leak that caused the server to crash, as well as improved documentation. The RLS is currently in use in production deployments in the GriPhyN LIGO and Earth System Grid projects.

### **Performance Evaluation of the Replica Location Service**

We have conducted an extensive performance evaluation of Version 2.0.9 of the Replica Location Service. A paper describing this implementation was accepted for publication in the High Performance Distributed Computing Conference (HPDC) to be held in Honolulu in June 2004. The submitted version of the performance study is available at:

<http://www.isi.edu/~annc/papers/chervenakhpdc13.pdf>.

### **Grid Service RLS Design**

We have continued to develop a specification for a Grid Service version of the RLS through the OREP (OGSA Data Replication Services) Working Group of the Global Grid Forum. The working group met at the GGF10 Meeting in Berlin in March 2004, where the latest version of the specification was discussed and the reference implementation for the Grid service was discussed.

### **Implementation of an OGSi RLS Grid Service**

We have built a reference implementation of the RLS Grid Service being standardized through the OREP group and conducted an initial performance evaluation of the service. A paper describing the implementation and its performance was presented at the GGF10 Data Area Workshop in Berlin in March. The paper is available at:

<http://www.isi.edu/~annc/papers/chervenakGGF10Data.pdf>.

### **Copy and Registration Service**

The Copy and Registration (CAR) Service was released through the Grid Technology Repository ([www.globus.org/gtr](http://www.globus.org/gtr)) in January 2004. The CAR Service provides a Grid service wrapper around the existing implementation of the RLS as well as the ability to perform file copy and delete operations and to have RLS servers reflect the creation or deletion of these files. Thus, the CAR Service represents a short-term solution for users who need a Grid service version of the RLS. The service is intended to be replaced eventually by the Grid Service RLS being designed through the OREP working group of the GGF.

### **Metadata Catalog Service**

We have continued the development of the Metadata Catalog Service. The latest implementation is based on the OGSA-DAI Database Access and Integration Service developed by the EPCC in Edinburgh. A paper describing the OGSA-DAI implementation as well as performance of that version and the earlier implementation based on Web services was accepted to the 16th International Conference on Scientific and Statistical Database Management (SSDBM04) that will take place in June 2004. The submitted version of this paper is available at: [http://www.isi.edu/~annc/papers/deelman\\_finall.pdf](http://www.isi.edu/~annc/papers/deelman_finall.pdf).

### **Batch-Aware Distributed File System**

This last quarter we have finished an initial prototype implementation and design of a new distributed file system designed for batch computing workloads. We published a description of this work in this years USENIX Networked Systems Design and Implementation conference. The contribution of this work is show how unique I/O sharing patterns within batch computing workloads can be exploited to increase throughput and improve failure handling. These benefits are achieved with a modified scheduler which controls storage decisions using mechanisms directly exported from the modified distributed file system.

### **Papers Published or in Progress**

Globus XIO: Enabling the Future of Grid Protocol Development, William Allcock, John Bresnahan, Rajkumar Kettimuthu, Joseph Link. Submitted to HPDC-13

### **Presentations Given**

“Overview of Globus Toolkit Data Management,” Globus World 2, January 2004.

“The Replica Location Service,” Globus World 2, January 2004.

“OGSA Data Replication Services Working Group,” Global Grid Forum 10 (GGF10), March 2004.

“Metadata Requirements for Replica Location Services and Metadata Catalog Services,” Metadata Research Group Birds of a Feather Session, GGF10, March 2004.

“A Replica Location Grid Service Implementation,” GGF10 Data Workshop, March 2004.

“Lessons Learned in Grid Networking” Panel Session at the 2<sup>nd</sup> International Grid Networking Workshop, CERN, March 15-17, 2004

“GridFTP Protocol Improvements” GridFTP WG co-chair GGF10, March 2004

“GridFTP Overview and how it could become Photonic Aware” OptIPuter Workshop, CalIT2, UCSD, March 2004

“Globus Toolkit Status and Futures”, Across Grids Conference, Nicosia, Cyprus, Jan 2004

“Overview of Globus Toolkit Data Management”, GlobusWORLD, San Francisco, CA, Jan 2004

“Globus XIO and GridFTP for Developers”, GlobusWORLD, San Francisco, CA, Jan 2004

“Firewalls in Grid Computing”, GlobusWORLD, San Francisco, CA, Jan 2004