

SciDAC DataGrid Middleware
A High-Performance Data Grid Toolkit:
Enabling Technology for Wide Area Data-Intensive Applications
Quarterly Report October 2002 thru December 2002

Accomplishments this Quarter:

Globus eXtensible IO (XIO) system design review completed.

The XIO code and documentation was distributed for review. Feedback was reviewed and incorporated into the design.

Globus XIO Framework, TCP, and File driver completed.

The framework, TCP, and File drivers have been completed. To aid in testing and transition, we have implemented the Globus IO API over XIO. Using this, with security disabled, we are currently passing the GridFTP test suite.

GridFTP server design started

The current GridFTP server is based on the Washington University FTP daemon (wuftpd). This provided several advantages: it allowed us to get the server out faster and it gave us credibility since it was a modification of a known entity. However, the wuftpd is extremely difficult to maintain and add features to. We also discovered a problem with licensing. For this reason, we are doing a complete re-implementation of the GridFTP server. The new server will be under the Globus Toolkit license and is being explicitly designed to be extensible. This will make the transition to an Open Grid Services Architecture (OGSA) much easier.

Code maintenance, development, hardening, collaborations and support.

Worked to continue developing and maintaining and hardening multiple software packages, including the NeST storage manager, the Bypass interposition agent and the Pluggable File System (PFS). Provided support for EDG users of Bypass and PFS and collaborations with SRM and ISI users of NeST. Also held discussions with members at CERN about using Bypass and PFS.

Building Migratory File Services for Grid Applications

Using the semantic information we had obtained from profiling the applications as described above, we enhanced a batch scheduling system with sophisticated storage management and workflow management software. We call this new system a *Migratory File Service* as it allows unmodified applications to migrate into remote environments transparently. Using interposition, the migratory file service recreates the namespace of the local environment. Using semantic information about the distinct types of I/O, the migratory file service eliminates batch and pipeline I/O from the home storage server thereby allowing much greater scalability and making the remote environment appear local in terms of performance as well.

The Ethernet Approach to Grid Computing

Developed a methodology and shell language for building robust systems in fault-prone environments with little or no failure detail.

Profiling Scientific Applications

Using interposition software and performance measurement counters on processors, we observed and measured the behavior of six different Grid application pipelines: BLAST, IBIS, CMS, HF, Nautilus and AMANDA. We found that I/O in these Grid applications could be split into three distinct types. *Endpoint* I/O is the unique input and output files for each application, *Batch* I/O is the shared dataset, which is read by multiple instances of each application, and *Pipeline* I/O is temporary data that is produced and consumed by different processes within the pipeline or even by different phases within an individual process.

Continued development, packaging, testing and deployment of the Replica Location Service

This quarter saw intensive continued testing of the Replica Location Service, including debugging and extensive functionality and performance testing. The RLS was packaged for release in Globus Toolkit 2.4 and in GT3 Alpha. The RLS was used by an increasing number of groups, including the LIGO physics application, the Earth Systems Grid and the Chimera system.

Turning the Replica Location Service into an Open Grid Services Architecture Service

Progress was made on creating OGSA services for the Replica Location Service. This was delayed somewhat by the lack of a C hosting environment in OGSA and by discussions within the Global Grid Forum about what the correct interfaces should be for the RLS services.

Began Re-design of the Metadata Catalog Service Based on OGSA-DAI

This quarter, we began redesigning the Metadata Catalog Service to be more general and scalable. The first step is to base our implementation on the OGSA-DAI database access and integration service from the UK EScience project. To date, we have deployed three versions of the OGSA DAI service, done extensive functionality and performance testing, and loaded metadata from the Earth System Grid project.

Plans for next quarter

Globus XIO Grid Security Infrastructure (GSI) driver and UDP driver complete

The last of the critical XIO drivers should be completed. The GridFTP test suite will pass with security enabled. The UDP driver can be used as the basis for the non-TCP protocol work.

Prototype of non-TCP XIO transport driver completed

We are examining RBUDP (UIC/Jason Leigh), Tsunami (IU/Steve Wallace), and SABUL (UIC/Bob Grossman) as options. We will have a prototype of at least one of these available for testing and evaluation.

Re-implemented GridFTP server should be able to move data with security features

Our server re-implementation should be well along at this point. The primary framework, the control channel, data channel, and basic security features should be in internal testing.

Resolving licensing issue for RLS

A licensing problem was found with the MySQL database currently used as the back end for RLS services. Unless this problem is resolved, it will be necessary for us to switch in the coming quarter to the postgresql database backend. This will require additional debugging and testing and possibly additional development.

Release of RLS in GT3.0 and possibly GT2.6

Once the database issues are resolved, we plan to release the RLS in stable Globus releases. This includes GT3.0 and also GT2.6, if there is such a release.

OGSA RLS Services

In this quarter, we will provide OGSA wrapper services around the existing RLS Local Replica Catalog and Replica Location Index services. These wrapper services will be deployed in the current OGSA Java hosting environment.

OGSA Copy and Registration Service

This quarter, we will also create an OGSA service that performs copy operations using the Reliable File Transfer Service and registers new replicas in the RLS. These operations will be reliable, with rollback of unsuccessful or partial operations.

Continued Redesign and Reimplementation of the Metadata Catalog Service

We will continue our redesign of the MCS in the coming quarter. In addition to using the OGSA-DAI service for individual metadata catalogs, we will begin the process of federating multiple metadata catalogs using a relaxed consistency model similar to that used in the Replica Location Service.

Papers Published or in Progress

“A Metadata Catalog Service for Data Intensive Applications”, Ann Chervenak, Ewa Deelman, Carl Kesselman, Laura Pearlman, Gurmeet Singh, in progress.

Douglas Thain, Todd Tannenbaum, and Miron Livny, "Condor and the Grid", in Fran Berman, Anthony J.G. Hey, Geoffrey Fox, editors, Grid Computing: Making The Global Infrastructure a Reality, John Wiley, 2003. ISBN: 0-470-85319-0

John Bent, Venkateshwaran Venkataramani, Nick Leroy, Alain Roy, Joseph Stanley, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau, and Miron Livny. “NeST - A Grid Enabled Storage Appliance,” In Jan Weglarz, Jarek Nabrzyski, Jennifer Schopf, and Maciej Stroinski, editors, Grid Resource Management, chapter 22. Kluwer Academic Publishers, 2003.

Douglas Thain, "The Ethernet Approach to Grid Computing" to appear in HPDC-2003,

This work has been accepted for publication into the 12th HPDC and is available here:
<http://www.cs.wisc.edu/~thain/workload-hpdc12.pdf>.

Presentations Given

January 14, 2003: Tutorial: GT2: Introductions and Basics at Globus World

January 14, 2003: Tutorial: GT2 Data Management at Globus World

January 14, 2003: Presentation on the Globus View of Data Management Architecture at GlobusWORLD.

January 16-17, 2003: Organized and led a Data Management Workshop at GlobusWORLD with attendance of approximately 50 participants from academia, government laboratories and industry.

January 17, 2003: Replica Location Service presentation at GlobusWORLD Data Management Workshop.

February 5, 2003: Presentations on the Replica Location Service, the Metadata Catalog Service and the Earth System Grid project at NASA Information Power Grid Meeting.

March 5, 2003: Birds of a Feather Session on OGSA Data Replication Services at Global Grid Forum in Tokyo, Japan.