SciDAC DataGrid Middleware
A High-Performance Data Grid Toolkit:
Enabling Technology for Wide Area Data-Intensive Applications
Quarterly Report April 2005 thru June 2005

Accomplishments this Quarter:

GT4.0.0 Released!
The latest stable release of the Globus Toolkit, GT4.0.0, came out on April 29th. This release represents a huge step forward across the board in terms of performance, reliability, ease of use, and documentation. This release contains several components supported by the SciDAC DataGrid Middleware project: GridFTP, the Reliable File Transfer (RFT) service, the Replica Location Service (RLS), and the eXtensible Input/Output (XIO) system.

Significant focus on user support
With the release of GT4 and in particular the new GridFTP server, we put a substantial amount of time into responding to support emails, configuration questions, etc.. The server is completely wire protocol compatible, but configuration and installation differ. Despite significantly improved documentation, "old hands" at GridFTP got caught by surprise. Whenever possible, responses included improvements to the online documentation.

Data Grid Architecture
We continue to be active with a number of groups, helping to define and refine the next generation of Data Grid architectures and design. A face-to-face design meeting was held in CERN with the European EGEE project, as well as a face-to-face and regular telecon calls with the Global Grid Forum (GGF) Open Grid Services Architecture (OGSA) Data Working Group. We also continue to be active in the Open Science Grid architecture and deployment. GridFTP is currently deployed and performing well and we are working on development of higher level transport services.

GridFTP Protocol V2.0 now a proposed standard
The new protocol which removes the "sender must connect" restriction as well as adding optional block by block "on the fly" check summing with resend to the data channel and new optional commands for the control channel.

GridFTP Data Storage Interfaces (DSI) continue to be implemented
The HPSS DSI is now successfully moving data. However, two issues have slowed our progress. The PIO API provided by IBM, and upon which our DSI depends, continues to have bugs. We have also been forced to add additional functionality that passes data between receiving DSIs so that the correct blocks are written to the correct tape drive.

We continue to work with the SCEC and Montage application communities to get GridFTP access to SRB data. A bug in the DSI was corrected to provide appropriate (non-delegated)

credentials to SRB. The SRB team does not view this work as a priority and this has slowed our progress considerably.

We are in design discussions with the Condor team to determine how to make the NeST DSI either the default, or at the very least a simple alternative to the POSIX DSI.

Replica Location Service Development

The Replica Location Service is included in the Globus Toolkit Version 4.0 release. We added two significant features to the RLS code during this quarter. One is the integration of RLS with the MDS monitoring infrastructure; RLS servers now provide information about themselves to MDS aggregators, which allow this information to be indexed and queried. Also, we added usage statistics to RLS. These are enabled by default and collect information such as the number of logical names, physical names and mappings registered in a catalog.

Aside from these new features, the main RLS code base is fairly stable, so we have mostly been improving documentation and working on minor bugs this quarter. We continue to support a growing number of applications that use the RLS in production, including the Laser Interferometer Gravitational Wave Observatory (LIGO) project, Earth System Grid, the US portions of the CMS and Atlas physics experiments, Nordugrid and others. This quarter, we have been helping the Earth System Grid project make the transition to the latest, most stable RLS version.

Also this quarter, we have begun a project to enable automatic configuration of RLS servers in large deployments. We assume that the Virtual Organization publishes a configuration file for RLS servers in the V.O. (for example, via an http server). We provide a mechanism that downloads the V.O. configuration file; possibly modifies it based on local requirements for access control and local customization such as back end database user name and password; and applies the configuration to a local RLS server. We provide scripts that merge V.O. and local configuration scripts and then signal the RLS server to apply these changes.

Finally, we have been working on a design for providing fine-grained authorization capabilities in RLS. We are seeking to provide this capability by building upon the Globus GT4 authentication infrastructure. There is an ongoing design discussion about how to implement this functionality using the GT4 custom authorization call-out.

WS-RF Data Replicator Service Released

This quarter, we completed the implementation of the Data Replicator Service (DRS). The first version was included as a Technical Preview Component for the Globus Toolkit Version 4.0 release. However, since that release, major additional work has been done on the service. This new version is available via CVS and will be included in the incremental 4.0.1 release. The goal of DRS is to provide a pull-based replication model to bring a set of requested files to a local site. The DRS interacts with the RLS to determine whether files exist locally, and if not, to find the locations where those files exist in the Grid. Then the DRS creates file transfer requests using the Globus Reliable File Transfer (RFT) service. The Data Replicator Service (DRS) is implemented in Java and complies with the WS-RF specifications.

Integration of Replica Location Service with POOL
We continued to work with the US CMS project on the integration of the Globus RLS with the POOL environment, which provides persistent object management for particle physics applications. This quarter, we successfully integrated this implementation into the latest POOL release.

Open Science Grid
We are leading the effort on defining data management architecture in the Blueprint document of the Open Science Grid activity.

OGSA Data Services Group
We continue to work with the OGSA Data Services Group in defining the data architecture document in the Global Grid Forum (GGF).

Studies of Interaction between the Network Storage (NeST) system and the Disk Resource Manager (DRM)
Reservation of storage space is a primary function of the Disk Resource Manager Interface and there is no major "competing" interface out there. However, it is rarely implemented because storage reservation is a difficult problem, but one that NeST directly addresses. For this reason we have implemented the storage reservation function of the DRM via NeST and have been running functionality and stability tests.

NeST Connection Manager Functionality
While management of the storage space is a major issue for large storage facilities, controlling the load on the system, i.e., making sure the system does not get swamped with requests and thus have all transfers fail is also an issue. We have added functionality to NeST that allows for connection management to throttle system access. This will allow a NeST administrator to limit the number of data transfers to/from a NeST, limit the size of transfers, etc. It does this through a concept of a connection reservation, allowing a user to reserve a number of connections in a manner similar to a NeST lot.

Plans for Next Quarter

1. Complete the HPSS DSI
2. Have the Montage and SCEC applications successfully using GridFTP to access SRB data in testing.
3. Add new functionality to GridFTP to allow SAML assertions to be sent over the control channel.
4. Continue general support for RLS, including bug fixes and improved documentation
5. Define a specification for a WS-RF wrapper around the RLS and request community feedback.

6. Complete design and begin implementation of fine-grained authorization capabilities for RLS.
7. Continued development of the Data Replicator Service. We will work with users in various communities who will experiment with the service and provide feedback.
8. Continue support for the RLS integration with the POOL infrastructure
9. Continue providing GT4 monitoring infrastructure that includes monitoring of RLS.
10. Continue interactions with Open Science Grid and OGSA Data Services groups.

Papers:

"Wide Area Data Replication for Scientific Collaborations", Ann Chervenak, Robert Schuler, Carl Kesselman, Scott Koranda, Brian Moe, submitted to the Grid2005 workshop, 2005.

"The Replica Location Service: Design and Experience", Ann Chervenak, Robert Schuler, Shishir Bharathi, Naveen Palavalli, Carl Kesselman, Ian Foster, Matei Reipeanu, Adriana Iamnitchi, being prepared for submission to IEEE Transactions on Storage.

"The Globus Striped GridFTP Framework and Server", W. Allcock, J. Bresnahan, R. Kettimuthu, M. Link, C. Dumitrescu, I. Raicu, I. Foster, SC'05, ACM Press, 2005.

Presentations:

"GridFTP: Easy Access to HPSS Data for Grid Applications" presented by Bill Allcock at the HPSS Users Meeting, Oakland, CA, Jun 2005.

"GridFTP vs the Firewall" presented by Bill Allcock at the Firewall Research Group meeting, Global Grid Forum 14, Chicago, IL, Jun 2005.