

Performance Characterization of GridFTP on 10 Gigabit Networks Using Hosts with 10Gb Network Interface Cards

J. E. H. Schwettmann, University of Oklahoma
R. Kettimuthu, MCS Division, Argonne National Laboratory

This project seeks to characterize the performance of Globus GridFTP transfers over both TCP and UDT protocols for hosts connected at 10Gb/s over a wide area network, and compare the performance of GridFTP using each. Standard disk-to-disk GridFTP transfers will be tested, as well as memory-to-memory transfers. Characterization will include detailed measurements regarding the speed and latency involved in the transfers, as well as resource utilization at the source and destination hosts.

What is Globus GridFTP ?

Globus GridFTP is an integral part of the open-source **Globus Toolkit**, a fundamental enabling technology for the “Grid”, and a central part of most Open Science Grid software installations.

GridFTP is part of the “Data Management” component of the Globus Toolkit, allowing FTP-like file transfers across grid-enabled resources. You may know it as **globus-url-copy**.

Recently, **UDT** (UDP Data Transfer) was added as an available protocol to GridFTP transfers, and though some early testing has been done, this project seeks to perform more extensive UDT-vs-TCP tests alongside others, and across a wider variety of testbeds

Previous Experiments

Bresnahan, et al. tested 1GB GridFTP transfers between several sites with a maximum connectivity of 1Gb/s, and found that GridFTP over UDT outperformed GridFTP over TCP by a **factor of 4** on two of the testbeds.

This early testing demonstrates that UDT has the potential to dramatically increase GridFTP throughput between Grid sites.

Ref: "UDT as an Alternative Transport Protocol for GridFTP", John Bresnahan, Michael Link, Rajkumar Kettimuthu and Ian Foster, Proceedings of the 7th International Workshop on Protocols for Future, Large-Scale and Diverse Network Transports (PFLDNeT 2009), Tokyo, Japan, May 2009

Automation with globus-transfer-test:

- **CPU & Memory Usage**—measured through top polling on source and destination machines, via globus-job-run
- **Parallel Streams**—Iterated or user-set
- **Protocols (TCP or UDT)**— iterated or user-set
- **TCP Buffer Size**—Calculated based on bandwidth and delay (ping via globus-job-run), or user-set
- **MTU size**—user-set (possibly iterated)
- **Transfer speed**—bytes vs wall time, and time/speed as reported by globus-url-copy
- **Disk-to-disk, Memory-to-memory, memory-to-disk**, and **disk-to-memory** tests are supported.
- Comparison with iperf, scp, and other common transfer and benchmark utilities to be implemented via globus-job-run.

Current Testbeds:

US ATLAS Tier2 Centers
OCHEP Tier2 Cluster
(University of Oklahoma)
Brookhaven National Laboratory
University of Michigan

Teragrid Resources
National Center for Supercomputing Applications (University of Illinois)
Indiana University
Purdue University
Oak Ridge National Laboratory
Louisiana State University
Pittsburgh Supercomputing Center
National Center for Atmospheric Research
Texas Advanced Computing Center (University of Texas)

globus-transfer-test (working name) is a utility being developed for automated testing of throughput and resource usage across any sets of Grid sites, and will provide valuable throughput optimization information for site administrators and end-users alike.

globus-transfer-test is being developed for potential inclusion into a future version of the Globus Toolkit.

Some Challenges to Characterization Testing

- 10Gb/s not often met in reality: many bottlenecks along the way.
- Bottlenecks can be difficult to isolate, netlogger (another new Globus Toolkit feature) makes this easier
- At 10Gb/s, fast disk access is crucial, and not always available.
- Maximum bandwidth can be difficult to determine without administrator input — this may be a challenge to future users of globus-transfer-test too.
- Test sites need to run the latest Globus Toolkit, with support for UDT, or support user sandboxes.

