

The nature of grid and distributed computing implies an inherent need for communication. Because of the various specializations of each site in a Grid environment and because some applications require use of more than one site, it is common for application users to have to move data from one site to another. The output of a large simulation computed at one site may need to be archived at another site and visualized for end users at a third site. This data is often very large, ranging from several hundred gigabytes to tens of terabytes. GridFTP is a well-known and robust protocol for fast data transfer on the Grid. It is based on the Internet FTP protocol, and it defines extensions for high performance operation and security. Implementations of it are widely deployed and used on well-connected Grid environments such as those of the TeraGrid because of its ability to scale to network speeds. Given resources, the GridFTP implementation provided by the Globus Toolkit can scale to network speeds and has been shown to deliver 27 Gb/s on 30 Gb/s links. The protocol is optimized to transfer large volumes of data commonly found in Grid applications. Datasets of sizes from hundreds of megabytes to terabytes and beyond can be transferred at close to network speeds by using GridFTP. The High Energy Physics community is basing their entire tiered data movement infrastructure for the Large Hadron Collider Computing Grid on GridFTP. The Laser Interferometer Gravitational Wave Observatory routinely moves 1TB a day during production runs. GridFTP is the recommended data transfer mechanism to maximize data transfer rates on the TeraGrid. Commonly used GridFTP clients include globus-url-copy, uberftp and Globus reliable file transfer service. Here, we describe a Globus XIO based client to GridFTP. Globus is a framework that presents a single standard posix-like open/close/read/write interface to many different protocol implementations. Protocol implementations can be incorporated into the XIO framework in the form of XIO drivers. A driver is responsible for manipulating and transporting the users data. XIO framework provides a well-defined interface for creating drivers. We developed a posix-like client for GridFTP by creating a GridFTP transport driver for Globus XIO. Such a client will greatly ease the addition of GridFTP support to third party programs, such as SRB, MPICH-G2 etc. Further, this client would provide an easier and familiar interface for applications to efficiently access remote files. In addition to the Open/Close/Read/Write interface, this client provides interfaces to utilize the features of GridFTP such as partial file transfer, parallel TCP streams, connection caching, setting TCP buffer size etc. Also, this client can be used to communicate with any FTP server. We compare the performance of this client with that of globus-url-copy on multiple endpoints in the TeraGrid infrastructure. We perform both memory-to-memory and disk-to-disk transfers and show that the performance of this posix-like client is comparable to that of globus-url-copy. Overall, this tool provides value to TeraGrid and other Grid users by making remote file access efficient and easy.