

FileTransfer

File Transfer

A common problem for many scientific applications is the replication of - often large - data sets (files) from one system to another. (For the generalized problem of transferring data sets from a source to multiple destinations, see DataDissemination.) Typically this requires *reliable transfer* (protection against transmission errors) such as provided by [TCP](#), typically *access control* based on some sort of authentication, and sometimes *confidentiality* against eavesdroppers, which can be provided by encryption. There are many protocols that can be used for file transfer, some of which are outlined here.

- FTP, the File Transfer Protocol, was one of the earliest protocols used on the ARPAnet and the Internet, and predates both TCP and IP. It supports simple file operations over a variety of operating systems and file abstractions, and has both a *text* and a *binary* mode. FTP uses separate [TCP](#) connections for control and data transfer.
- HTTP, the Hypertext Transfer Protocol, is the basic protocol used by the World Wide Web. It is quite efficient for transferring files, but is typically used to transfer *from* a server to a client only.
- RCP, the Berkeley Remote Copy Protocol, is a convenient protocol for transferring files between Unix systems, but lacks real security beyond address-based authentication and clear-text passwords. Therefore it has mostly fallen out of use.
- SCP is a file-transfer application of the [SSH](#) protocol. It provides various modern methods of authentication and encryption, but its current implementations have some *performance limitations* over "[long fat networks](#)" that are addressed under the [SSH](#) topic.
- BitTorrent is an example of a *peer-to-peer file-sharing* protocol. It employs local control mechanisms to optimize the global problem of replicating a large file to many recipients, by allowing peers to share partial copies as they receive them.
- [VFER](#) is a tool for high-performance data transfer developed at Internet2. It is layered on [UDP](#) and implements its own delay-based rate control algorithm in user-space, which is designed to be "TCP friendly". Its security is based on [SSH](#).
- [UDT](#) is another UDP-based bulk transfer protocol, optimized for high-capacity (1 Gb/s and above) wide-area network paths. It has been used in the winning entry at the Supercomputing'06 Bandwidth Challenge.

Several high-performance file transfer protocols are used in the Grid community. The "comparative evaluation..." paper in the references compares *FOBS*, *RBUDP*, *UDT*, and *bbFTP*. Other protocols include [GridFTP](#), [Tsunami](#) and [FDT](#). The eVLAB community uses file transfer tools from the Mark5 software suite: [File2Net](#) and [Net2File](#). The ESnet "[Fasterdata](#)" knowledge base has a very nice section on [Data Transfer Tools](#), providing both general background information and information about several specific tools. Another useful document is Harry Mangalam's [How to transfer large amounts of data via network](#) a nicely written general introduction to the problem of moving data with many usage examples of specialized tools, including performance numbers and tuning hints.

Network File Systems

Another possibility of exchanging files over the network involves *networked file systems*, which make remote files transparently accessible in a local system's normal file namespace. Examples for such file systems are:

- [NFS](#), the Network File System, was initially developed by Sun and is widely utilized on Unix-like systems. Very recently, NFS 4.1 added support for pNFS (parallel NFS), where data access can be striped over multiple data servers.
- [AFS](#), the Andrew File System from CMU, evolved into [DFS](#) (Distributed File System)
- [SMB](#) (Server Message Block) or [CIFS](#) (Common Internet File System) is the standard protocol for connecting to "network shares" (remote file systems) in the Windows world.
- [GPFS](#) (General Parallel File System) is a high-performance scalable network file system by IBM.
- [Lustre](#) is an open-source file systems for high-performance clusters, distributed by Sun.

References

- *A Comparative Evaluation of High-Performance File Transfer Systems for Data-intensive Grid Applications*, C. Anglano, M. Canonico, Proc. 13th IEEE Int. Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises (WET ICE'04) [PDF](#), [requires IEEE account](#)
- [fasterdata](#), ESnet [<http://fasterdata.es.net/fasterdata/>] Network Performance Knowledge Base, February 2011 (last checked)
- [How to transfer large amounts of data via network](#), Harry Mangalam, Web page, April 16 2016 (version accessed)

– Main.TobyRodwell - 2005-02-28

– Main.SimonLeinen - 2005-06-26 - 2015-04-22