Towards the Development of a Pipelined Data Movement with Globus Online

Christopher Paolini
San Diego State University
5500 Campanile Drive
San Diego, CA 92182
paolini@mail.sdsu.edu

Mary P. Thomas
San Diego State University
5500 Campanile Drive
San Diego, CA 92182
mthomas@mail.sdsu.edu

Madhukar Yedulapuram
San Diego State University
5500 Campanile Drive
San Diego, CA 92182
madhukary@hotmail.com

Raj Kettimuthu
Argonne National Laboratory
9700 S Cass Avenue
Argonne, IL 60439
kettimut@mcs.anl.gov

ABSTRACT

Big data has become an inevitable part of science. Data generated from experiments run at research facilities and simulations performed at computing centers has to be transferred to other locations for analysis, visualization, and archival purposes. For example, XSEDE migrated its network backbone infrastructure to the 100G Internet2 Network, which offers multiple paths between 22 connector sites that service university campus and national lab endpoints. Similarily, San Diego State University has installed two networks: the CSRCNet network (a direct 10Gb/s connection to SDSC resources); and an NSF Campus Cyberinfrastructure Network Infrastructure and Engineering Program (CC-NIE) Science DMZ which will host a separate network optimized for large scientific data transfer to other remote resources [1]. The time required to transfer between such sites is a function of approach, transfer mechanisms used, file size, and the bandwidth available on the network. Globus Online makes the task of data transfer easier for users by allowing them to outsource the task of data transfer to the hosted Globus Online service [2]. Globus Online uses GridFTP to get high performance, monitors the transfers, and restarts transfers automatically in case of any transient failures. The transfers are still limited by the bottleneck network bandwidth between the source and the destination. In cases where the standard network path between the source and destination is slow, it may be possible to find a better alternative path using an intermediate host. Although GridFTP can use UDT, an application-level protocol that uses UDP for data transport, TCP is the default transport protocol, which can be detrimental to the transfer performance of scientific data over high-delay networks. On such networks, congestion will lead to packet loss, which causes TCP to reduce, significantly, its sending rate. The sending rate will increase if packet loss subsides, but will be continually reduced if packet loss resumes. In fact, a packet loss of 1% can result in a 100X decrease in TCP sending rate. To improve transfer times, users of XSEDE resources who encounter significant performance reductions while transferring datasets between their campus endpoint and an XSEDE system may choose an alternative path with a lower packet loss rate that traverses different Internet2 connector sites. Such an alternative path may be defined by specifying a mix of intermediary XSEDE systems that lie on the alternative path. Furthermore, some users may have high-speed, low-latency dedicated connections to a particular, geographically near XSEDE resource who may want GridFTP transfers from other, distant XSEDE systems to pass through the dedicated connection. By specifying the geographically near XSEDE resource as an intermediary host, one can use GridFTP to transfer a dataset from a distant XSEDE system to the near XSEDE system, and finally from the near XSEDE system to a campus endpoint, forcing data to traverse the user's preferred, dedicated network. In this paper we describe our initial reports to improve the performance of Globus Online transfers by pipelining data transfer through an intermediate host. We also describe the applicability of this solution in the context of a real world application involving coastal ocean model simulations [3], run on XSEDE resources.

1. REFERENCES

[1] R. P. J. E. Castillo. C. P. Paolini, "Campus Cyberinfrastructure - Network Infrastructure and Engineering Program (CC-NIE)(nsf12541)," 2012.

[2] D. Sulakhe, R. Kettimuthu, and D. Utpal, "Sequencing Centers Using Globus Online : A Case Study," in *Proc. Intl. Conf on Computing in High Energy and Nuclear Physics (CHEP)*, pp. 1–6, 2012.

[3] M. P. Thomas and J. E. Castillo, "Parallelization of the 3D Unified Curvilinear Coastal Ocean Model : Initial Results," in *12th International Conference on Computational Science and Its Applications*, (Salvador de Bahia, Brazil), 2012.