Providing Real-time Data Transfer Functionality Without Bandwidth Reservations

NAME: Gagan Agrawal, Rajkumar Kettimuthu, Ian Foster
AFFILIATION: Ohio State, Argonne National Laboratory
EMAIL: agrawal@cse.ohio-state.edu
EXPERTISE: Scheduling, Middleware Systems, Big Data Applications, Analytical & data-driven models,

Many (DOE) science applications are requiring real-time wide-area data transfers. Such transfers are not feasible today on most networks, though need (and mechanisms) for supporting reservations has been discussed in the literature, and some of the newer networks support bandwidth reservations. We argue for a different approach, which involves scheduling in a way that near real-time performance is achieved by a certain set of transfers.

MOTIVATION

The need to ensure timely completion of data transfers arises in multiple science communities. Many relevant science cases involve an instrument that produces data: as data volumes (and thus computational requirements) increase at light sources, fusion reactors, and other experimental facilities, sufficient computing power is no longer available locally. Thus, researchers need to depend on remote computing facilities. Data transfer and analysis must happen in a timely manner in order to check results, adjust the experimental setup, and maximize the use of experimental facilities.

We have recently extensively studied scheduling of wide area transfers. Based on the observation that concurrency (no. of partial files a given file is broken into) and/or use of parallel TCP streams can be used to control the bandwidth used by a given transfer, we have developed nuanced scheduling algorithms. Among these algorithms, one particular algorithm focuses on prioritizing real-time transfers, while also not negatively impacting other transfers. The results of applying this algorithm on real logs for transfers between Stampede and several other XSEDE sites have been impressive.

POTENTIAL BENEFITS

1) Reduced investment in both initial hardware and network operations by alleviating the need for bandwidth reservations
2) Improved support for large-scale DOE data-sciences.

KEY DOE NETWORK/SCHEDULING CHALLENGES

1. What control knobs would exist in future networks so that bandwidth allocation can be controlled at the application/transport later?
2. How do we build models that relate achieved bandwidth to the values of the control parameters? How does one account for external load, which are the transfers originating outside of our control? Can the network provide information about external load to the user via simple and well-defined APIs?
3. How does one scale the scheduling techniques (and supporting middleware) to work with a large number of end-points?