

Project Summary / Abstract

BigData Express—Toward Schedulable, Predictable, and High-Performance Data Transfer

Problem: In DOE research communities, the emergence of distributed, extreme-scale science applications is generating significant challenges regarding data transfer. We believe that the data transfer challenges of the extreme-scale era are characterized by two relevant dimensions:

- High-performance challenges. The DOE is working toward deploying terabit networks in support of extreme-scale science applications. Ideally, high-performance data transfer will reach terabit/s throughput to make full use of the underlying networks.
- Time-constraint challenges. Scientific applications typically have explicit or implicit time constraints on data transfer. Providing real-time and deadline-bound data transfer is a challenging task in the extreme-scale era.

Although significant improvements have been made in the area of bulk data transfer, currently available data transfer tools and services will not be able to successfully meet these challenges, for the following reasons:

- Existing data transfer tools and services lack a data-transfer-centric approach to seamlessly and effectively integrating and coordinating the various entities in an end-to-end data transfer loop.
- Existing data transfer tools and services lack effective mechanisms to minimize cross-interference between data transfers.
- Existing data transfer tools and services are oblivious to user (or user application) requirements (e.g., deadlines and QoS requirements).
- Inefficiencies arise when existing data transfer tools are run on DTNs.

These are common and fundamental problems for bulk data transfer in the extreme-scale era. In this proposal, we seek to address these problems.

Proposed Solution: We propose to research, develop, and implement “*BigData Express—Toward Schedulable, Predictable, and High-Performance Data Transfer.*” The proposed BigData Express is a middleware data transfer service that would have the following features:

- A data-transfer-centric architecture to seamlessly integrate and effectively coordinate the various resources in an end-to-end data transfer loop.
- Employment of SDN and SDS to improve network and storage I/O performance for BigData Express.
- A time-constraint-based scheduler to schedule data transfer tasks.
- An admission control mechanism to provide guaranteed resources for admitted data transfer tasks.
- A rate control mechanism to improve data transfer schedulability and reduce cross-interference between data transfers.

BigData Express seeks to provide a *schedulable, predictable, and high-performance* data transfer service for DOE’s large-scale science computing facilities (LCF, NERSC, and US-LHC computing facilities, among others) and their collaborators.

Proposed Team: The proposed BigData Express is a joint medium-sized project between FNAL and ORNL. In addition, ESnet, as an unfunded project partner, will provide the underlying SDN-based WAN services required for a successful project. The principal investigator (PI), Dr. Wenji Wu, is the lead scientist for network research at Fermi National Accelerator Laboratory