

**Collaborative Research: SDCI Net:
Policy-driven Large Scale Data Access Framework
with Light-weight Performance Monitoring and Estimation
(Adaptive Data Access and Policy-driven Transfers - ADAPT)**

Award Numbers: 1127101 (USC/ISI), 1127039 (LBNL)

Outcome Report

April 15, 2015

Principal Investigators

Ann Chervenak¹, Alex Sim²

Project member: Robert Schuler¹, Junmin Gu²

Project Website: <http://sdm.lbl.gov/adapt/>

Project email: adapt@hpcrd.lbl.gov

NSF Program Officer: Kevin L. Thompson

¹ University of Southern California, Information Sciences Institute

² Lawrence Berkeley National Laboratory

The overall goal of the Adaptive Data Access and Policy-driven Transfer (ADAPT) project is to improve the performance of large data transfers between two sites connected by the Internet. Such large data transfers are often performed by distributed scientific collaborations. For example, a scientist may want to analyze a data set that is stored at one site, but the analysis itself needs to take place at a second site that has the required computational resources, such as a supercomputer or a large computing cluster. In such cases, the data set must first be copied to the computational site before the analysis can begin.

These large data transfers are often inefficient using current data transfer technologies. Often the data set is stored in many (hundreds or thousands of) files that must all be sent to the target site. If the resources are not carefully managed at the source and destination sites and on the network between them, it is easy to overwhelm those resources. For example, the source site may try to send the data faster than the network or the destination site can transfer or store the data. This results in poor transfer performance, in terms of *latency* (how long it takes for data to arrive) or *throughput* (the rate of transferring data).

The ADAPT project studied two techniques to improve the performance of these large data transfers. First, the project studied how to manage the resources available for transferring data more effectively. This *resource management* is based on the knowledge of the distributed scientific collaboration regarding how it wants to manage its data transfers (for example, that transfers between certain sites should get priority over other transfers) as well as any limitations known by individual sites (for example, that the site's storage systems can only transmit or receive data at a certain rate). Such preferences and limitations can be considered the *policies* of the underlying systems. The ADAPT project implemented a Policy Service that is responsible for allocating system resources based on these preferences.

Second, the ADAPT project improved the performance of individual data transfers over time by adapting their transfer characteristics based on how well the transfers are currently performing. Characteristics of a transfer might include how many parallel streams are used to transfer one file. The ADAPT team enhanced existing data transfer software with additional capabilities to check recent performance and adapt transfer characteristics accordingly. If recent performance was good, then the software will increase its use of streams to achieve even greater bandwidth for the transfer. If recent performance was poor, then the software will decrease its use of streams to avoid overwhelming the available resources.

The project developed open source software that is freely available, including (1) the Policy Service that is responsible for allocating resources based on preferences of the sites and the distributed collaboration, and (2) the enhanced data transfer software that is responsible for adapting transfer characteristics based on recent performance. Figure 1 shows how these software components interact during data transfers, with each transfer client contacting the Policy Service to get an allocation of resources and then adapting the characteristics of the transfers that it performs within that allocation.

The ADAPT team ran experiments with these two techniques and demonstrated that these techniques are very effective in improving the overall performance for large data transfer operations between two sites. The techniques work particularly well in environments where the available resources, either at the source and destination sites or on the networks between them, are limited. Such limitations frequently occur on real systems, for example, when site resources are shared by multiple users who may all be performing competing data transfers and computations. Likewise, Internet bandwidth may be limited between two sites because there are multiple transfers competing for the available network resources.

The ADAPT techniques are successful in efficiently allocating resources and using transfer adaptation to avoid overwhelming the available resources. Figure 2 shows the results of experiment. In this scenario, a scientist is transferring a large data set from a site in Korea over the Internet to a site in Oakland, CA, where an analysis will be performed on the data set. The red line shows the throughput (data transfer per second) over time using the adaptive techniques compared to conventional data transfers, shown on the

black line. The ADAPT techniques show a 20% improvement in the overall time to complete the large transfer between the sites.

The results of the ADAPT project will have broad impact on many scientific domains that need to transfer large amounts of data efficiently between sites over the Internet using limited, shared resources. In addition, these techniques will be relevant to other areas, including Software Defined Networking, where the resource allocation and adaptive transfer techniques could help improve the routing of network traffic.

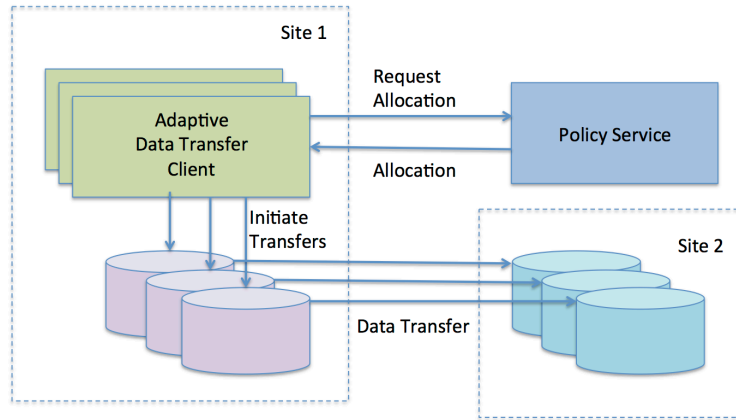


Figure 1: Service interactions for bulk remote data access use cases

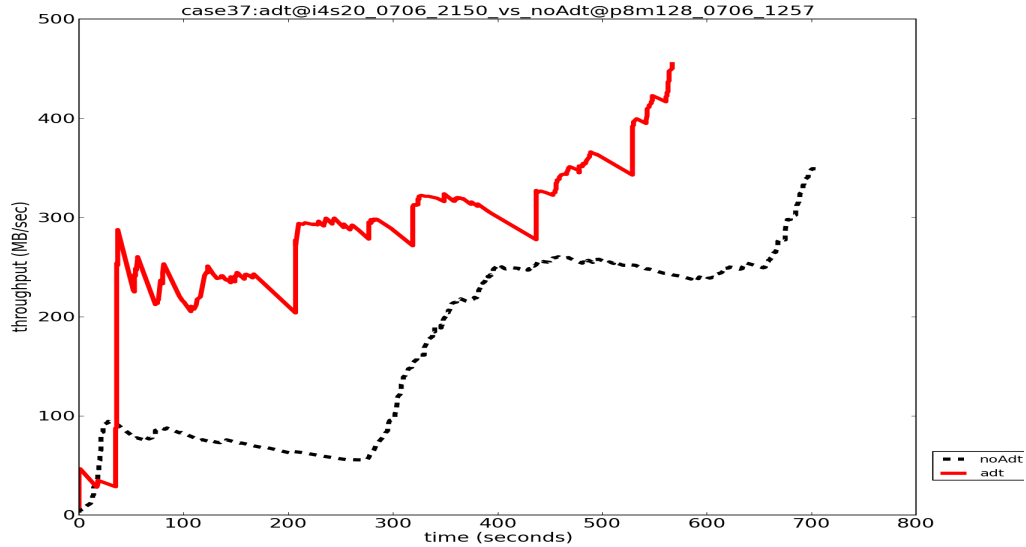


Figure 2. Comparison of adaptive vs. non-adaptive data transfer performance for a 260 GByte data set from NISN in Korea to NERSC in Oakland, CA. The red line shows the aggregate throughput in MBytes/second when using both ADAPT resource allocation and adaptive data transfers; the black line shows the throughput transfer without the ADAPT techniques. The ADAPT transfers complete approximately 20% faster.