

# DATA STORAGE

## PROVIDING DATA ON TAPE AND DISK

Fermilab provides long-term custodial storage of tens of petabytes of scientific data. Scientists on-site can directly access files on tape through Fermilab's Enstore Mass Storage System. From on- or off-site, high-rate access to files is through dCache, a disk-cache front-end to the tape storage that can also be used alone. Fermilab has served data at 45 terabyte/hr on-site and 2 terabyte/hr off-site.

dCache, developed jointly by DESY, NDGF and Fermilab, provides high-performance disk storage as a high-speed front end to a tape or standalone system. It supports a variety of transport protocols and authorization schemes. The Open Science Grid supports and distributes dCache as part of its Virtual Data Toolkit.



The DZero experiment at Fermilab, an Open Science Grid user, typically submits 60,000-100,000 simulation jobs per week. OSG worked with member institutions to allow DZero to use opportunistic storage, that is, idle storage on shared machines, at several sites. With allocations of up to 1 terabyte at processing sites, DZero increased its job success rate from roughly 30% to upwards of 85%.

Fermilab supports dCache installations in North America and beyond for the LHC experiments at CERN and Open Science Grid sites, including one of the largest in the world at its US CMS Tier-1 site.

TEVATRON EXPERIMENTS CURRENTLY STORE MORE THAN 3 PETABYTES (3 BILLION MEGABYTES) PER YEAR.



# NETWORKING

## BUILDING AND MONITORING NETWORKS AND OPTIMIZING DATA FLOW

### BUILDING A NETWORK

Built and operated jointly with Argonne National Laboratory, Fermilab's Metropolitan Area Network (MAN) expects to reach over 600 gigabytes per second potential capacity, with 80 gigabytes per second active now. This network, and redundant equipment on-site, have kept Fermilab free from internet outages for over two years.

### MONITORING NETWORKS

Fermilab contributes to the development of the perfSONAR network monitoring middleware that is designed to facilitate troubleshooting across multiple network management domains.

Fermilab has deployed perfSONAR to monitor two major Large Hadron Collider-related, end-to-end WAN infrastructures: the LHC Optical Private Network (LHCOPN) between Fermilab and CERN in Switzerland, and a heterogeneous set of network links between Fermilab and about a dozen U.S. university sites for the CMS experiment at CERN.

Fermilab leads the new Network Weather and Performance E-Center project that will establish a facility where users can obtain information on desired end-to-end network paths. PerfSONAR services will underpin its data collection and performance measurement.

Fermilab participates in the End Site Control Plane System project to enable end sites, e.g., universities, to establish a data circuit to a collaborating site on the general U.S. research and education network infrastructure.

### OPTIMIZING DATA FLOW

The joint Fermilab-Caltech Lambda Station project enables the dynamic rerouting of designated traffic through site LAN infrastructure onto so-called "high-impact", wide-area networks. For the large, sustained flows of high-energy physics data, Lambda Station works in an application-independent reactive mode, in which its action is triggered by changes in existing network traffic.

HIGH-CAPACITY NETWORK CAN MOVE 100 PETABYTES OF DATA PER YEAR, AND ENABLE REAL-TIME COLLABORATION.

