Global Terascale Data Management for Legacy Applications

Logistical netCDF (netCDF/L) is an enhanced version of the popular network Common Data Form (netCDF) that enables users to easily store, replicate, and control the location of their large netCDF files on the global Logistical Network (30+TB and 400+nodes). With netCDF/L, application communities, such as the Terascale Supernova Initiative (TSI), achieve ease of use, high performance and flexible control of their "data logistics" by using metadata files called exNodes and a directory service for managing them – the Logistical Distribution Network (LoDN, "low down").

- **Familiar netCDF interface**
  NetCDF/L meets the all netCDF acceptance tests and works with legacy applications unmodified.
- **Single upload, multiple download**
  Logistical Networking combines the collaborative power of high performance networking with the ease of use and control of a shared global file system.
- **Point-to-multipoint (over time)**
  In the Logistical Network, data can be written to depots at one location initially, but later moved/replicated to other locations, depending on changing needs and conditions. LoDN makes such "data logistics" transparent to the application.
- **High Performance Downloads**
  NetCDF/L uses high performance algorithms to accelerate data transfers, routinely achieving hundreds of Mbps across research networks, like Abilene.
- **Selective Downloads**
  Downloading the entire content of a large data set from the network is sometimes unnecessary or undesirable (e.g. due to bandwidth or storage limitations). With netCDF/L, users can read selectively only the parts of the file that are needed.

netCDF

NetCDF is a library for manipulating locally stored datasets as a structured collection of variables that can be read and written independently of one another. In the figure at the right, X, Y, and Z represent 3 different netCDF file segments associated with 3 different variables. The figure below illustrates how the TSI community uses netCDF/L to leverage this capability on globally stored datasets.

NetCDF/L and TSI

In collaborative environments, not every variable is always required by every participant. This provides a basis for differential movement of data based on the application requirements, making the best possible use of available storage and network resources. TSI simulations that run on the Cray X-1 at Oak Ridge write their output to the US Logistical Network in netCDF format using netCDF/L. Different TSI sites can then use netCDF/L to read/download just the part of the simulation data they require.

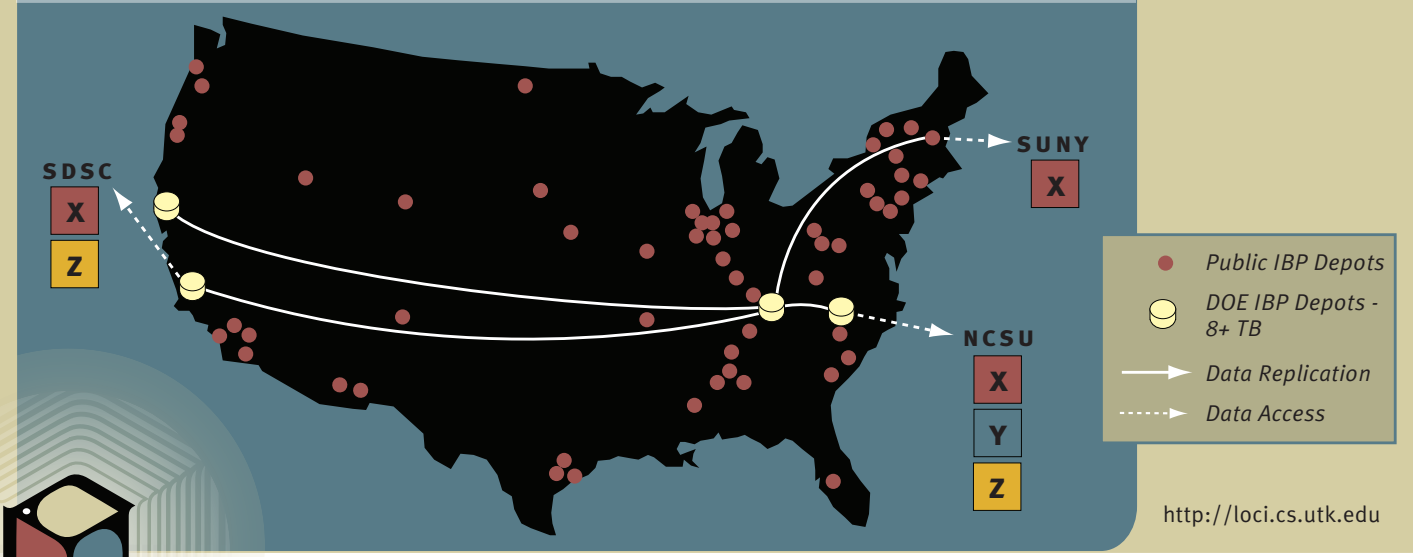

http://loci.cs.utk.edu

LoCI
Logistical Computing & Internetworking Lab
Department of Computer Science
University of Tennessee

NSF DOE Internet2 PlanetLab CITR UT