U.S. DEPARTMENT OF ENERGY | Office of Science

Data Storage and Data Transfer

Dan Hitchcock

Acting Associate Director, Advanced Scientific Computing Research

Office of Science Organization

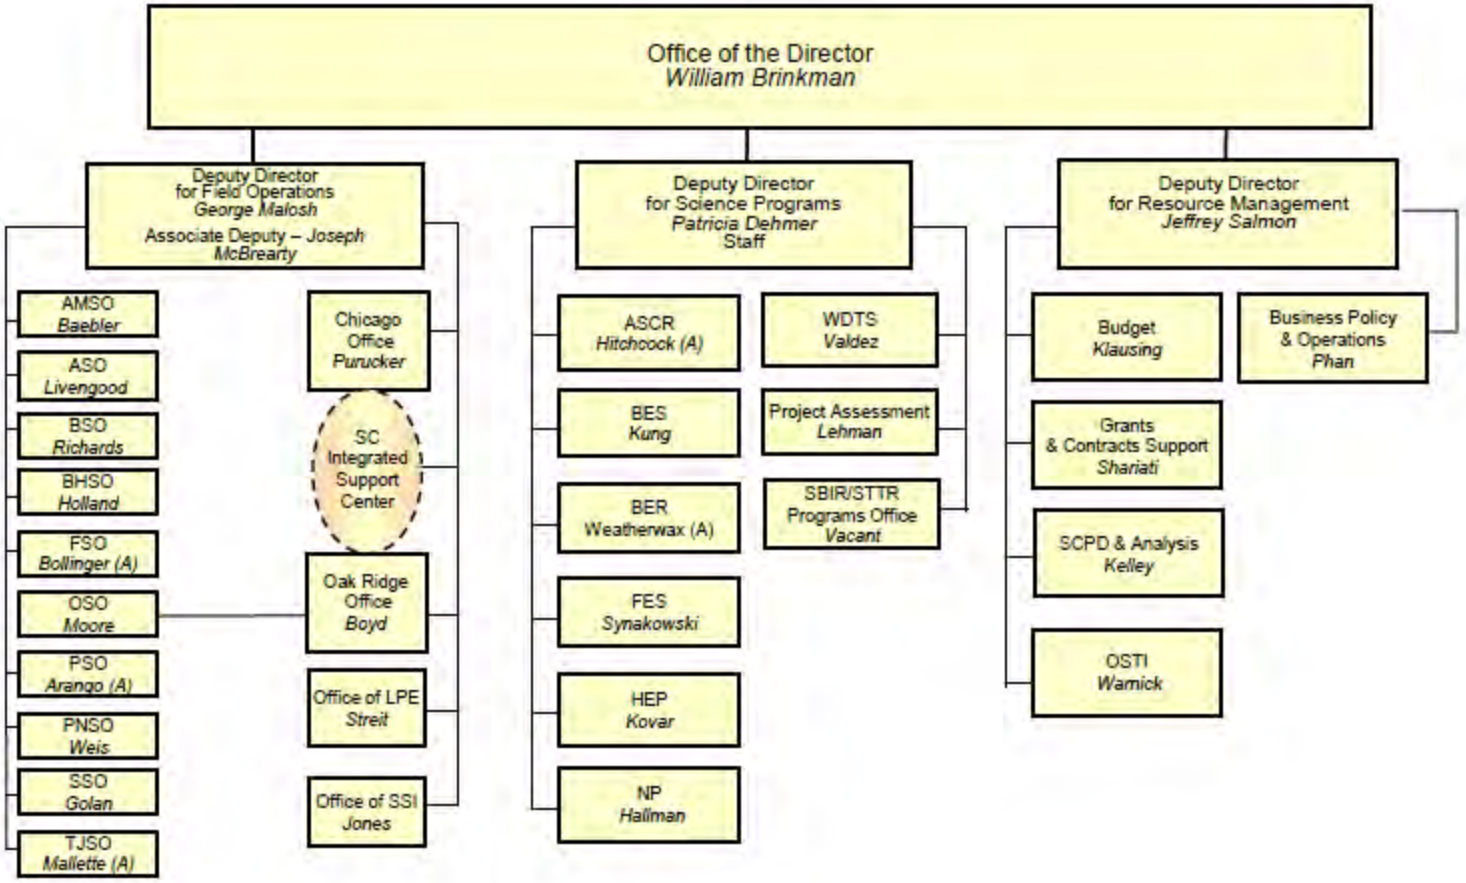

Data to Enable Scientific Discovery

- **DOE SC Investments in Data Driven Science support SC and DOE missions:**
  - **Advanced Scientific Computing Research: Leadership Computers and NERSC**
  - **Basic Energy Sciences: Light Sources, Neutron Sources,…**
  - **Biological and Environmental Research: Climate, Proteomics, …**
  - **Fusion Energy Sciences: Experiments and Simulation**
  - **High Energy Physics: LHC,…**
  - **Nuclear Physics: TJNAF, RHIC,…**
- **Data and Science closely coupled**
- **End to End perspective**

Requirements Gathering

- **ESnet and NERSC separately survey each SC program office every 2-3 years (two per year) to determine program's computing and networking requirements**
- **Workshops jointly organized by ASCR and Program Office**
- **Program Office chooses representatives**
- **Findings from the NERSC-BES Requirements workshop, February, 2010**
  - **Several large BES experiments coming online that will have large data requirements. For example, the Coherent X-ray Imaging instrument at SLAC is expected to generate 1.7 PB of data in 2012 and 10.4 PB in 2015**
  - **Other than these large experiments, all BES projects presented at the workshop had data needs that are currently being met and are projected to being met in the future based on our current data growth projections.**

BES ESnet Requirements Workshop Highlights

- **Instrument upgrades are driving dramatic changes in science process for BES facility users**
  - **Current paradigm is to take data home on portable USB hard drives**
  - **Current paradigm tractable because data sets are small, and end sites have poor network performance characteristics**
  - **Instrument and facility upgrades are increasing data set sizes by up to 100x in the next few years (individual beamlines at some facilities, significant portion of NSLS-II at BNL in 2014)**
  - **New paradigm: data sets too large, not portable, must be analyzed at facility and/or transferred to home institution**
- **Remote control will become more important as experiments become increasingly automated**
  - **Science Data Network (SDN) has a role to play here (we have seen this coming)**
  - **Need to vet operational models with facilities**
- **BES facilities and users are going to need a lot of help from ESnet, site networks, and tools providers – ESnet already planning pilot projects (ALS, NSLS)**

Other BES-ESnet Workshop Highlights

- **Multiple scientists expressed a need for remote visualization – this will increase as data sets increase in size (data too large to transfer home → analyze data at remote institution → remote visualization need)**
- **Significant knowledge gap in science community regarding appropriate tools, proper setup for data transfer, etc**
  - **http://fasterdata.es.net is helping, but there is more to do**
  - **ESnet working with facilities (network tuning talk at workshop generated significant interest from attendees)**
- **Cybersecurity is a major obstacle**
  - **SSH/SCP are used for data transfer (installed by default and approved by cybersecurity) and are the worst performers for data transfer**
  - **Appropriate tools (e.g. GridFTP) need to be deployable in order to be useful – if cybersecurity prevents their deployment, they do not exist as useful scientific tools**
- **Portals and other data services (e.g. NERSC Science Gateways) were of interest to several participants**

ESnet Backbone – Connecting Facilities

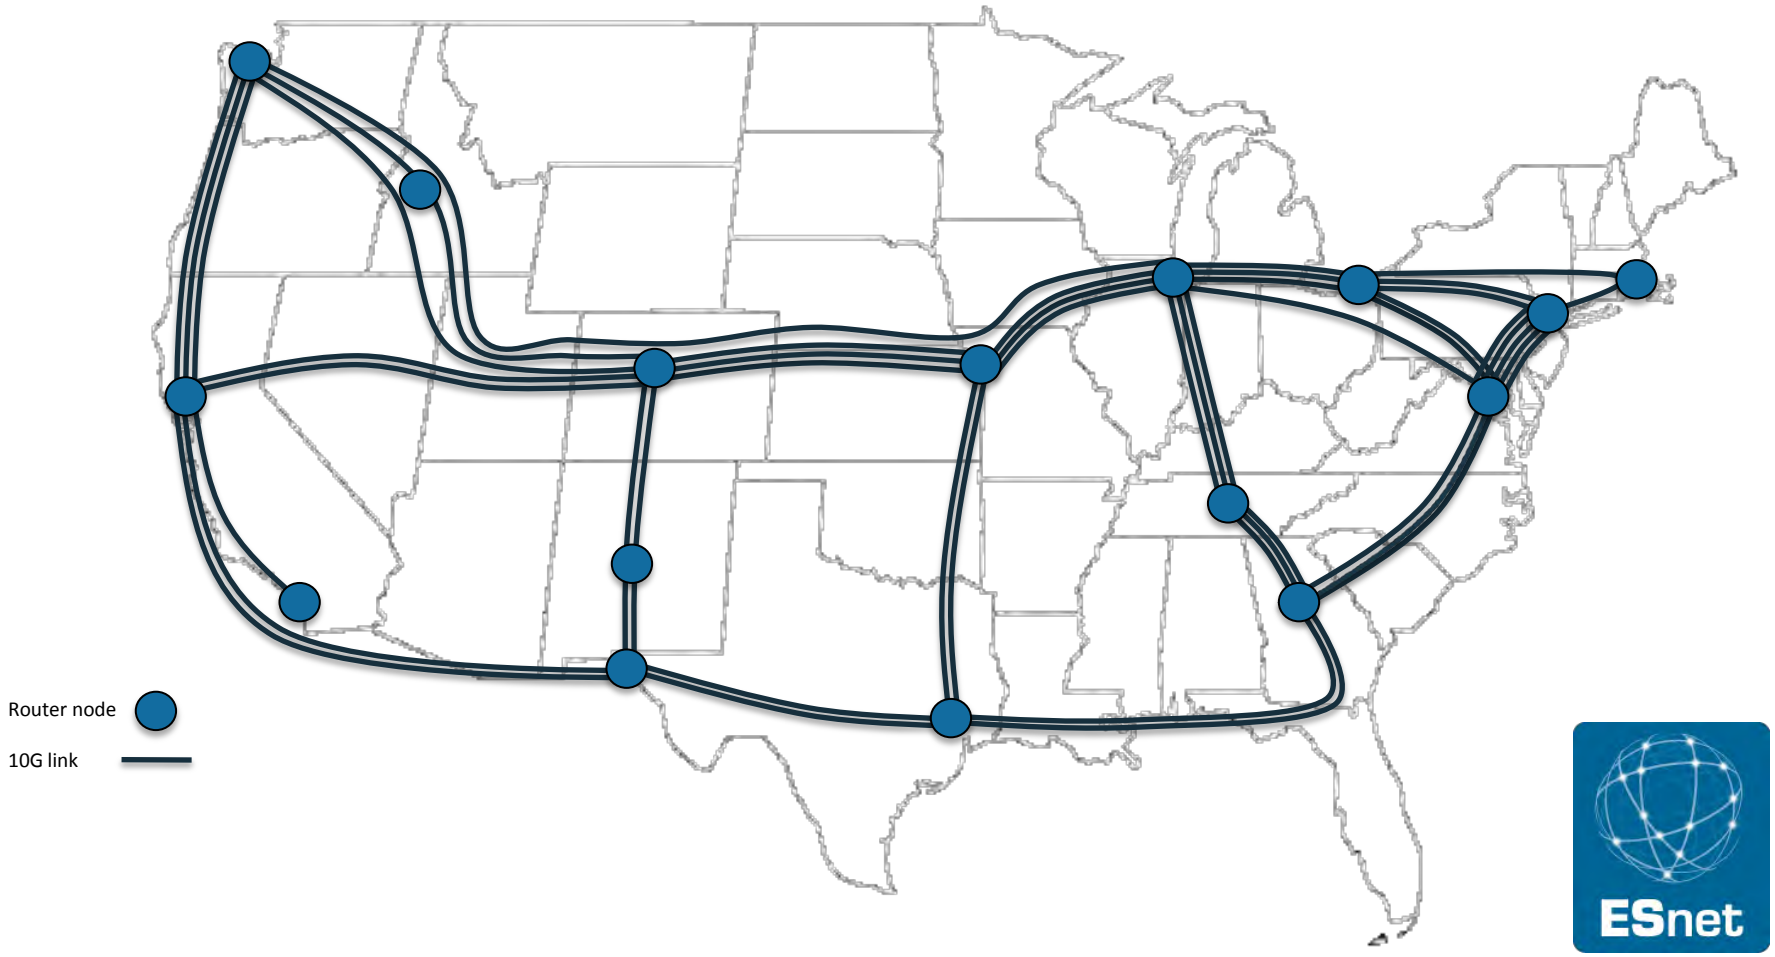

ESnet

Data Transfer Tools

The Center for Enabling Distributed Petascale Science (CEDPS) provides secure, high-bandwidth "fire-and-forget" file transfer services

- **CEDPS project team: ANL, LBNL, USC/ISI, FNAL, UW-Madison**
- Serving scientists and resource providers from: LBNL / NERSC, ORNL, ANL, GFDL, TeraGrid, Europe, Australia

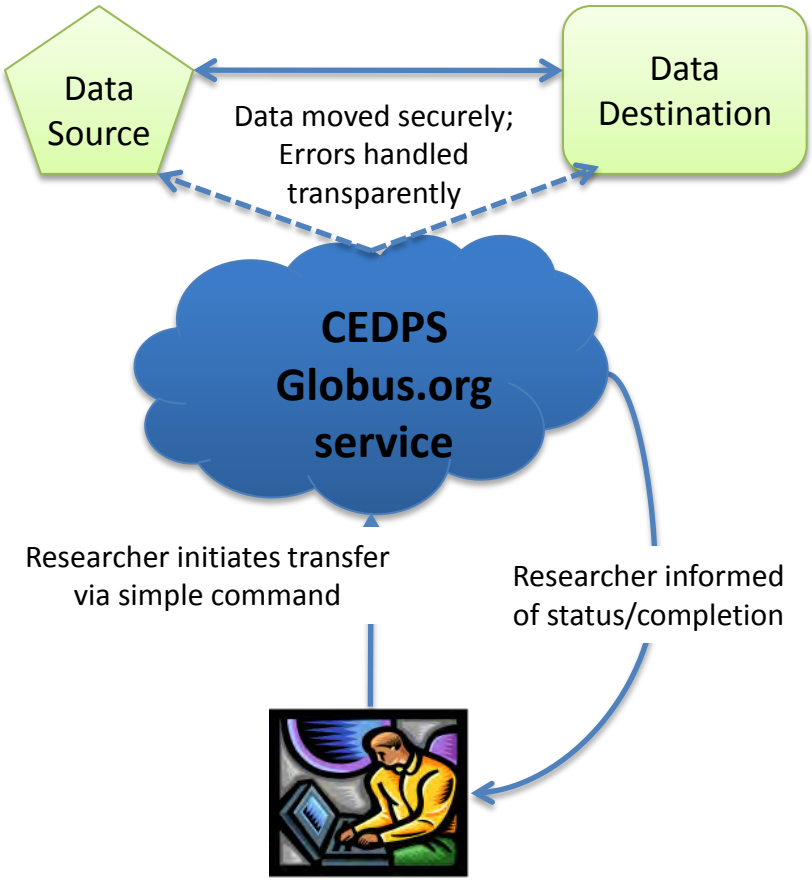

Accomplishments

Transferred without human intervention, automatically recovering from many network and systems failures:

- **STAR: 300GB of data between BNL and NERSC/PDSF**
- **SCEC: 600,000 4MB files from TeraGrid resources to OSG**
- **CEDPS: 20TB among 12 sites (including DOE labs, TeraGrid, Australia) in under 15 hours**

Powered by:

Uniformity of Interface ➔ Compatibility of SRMs

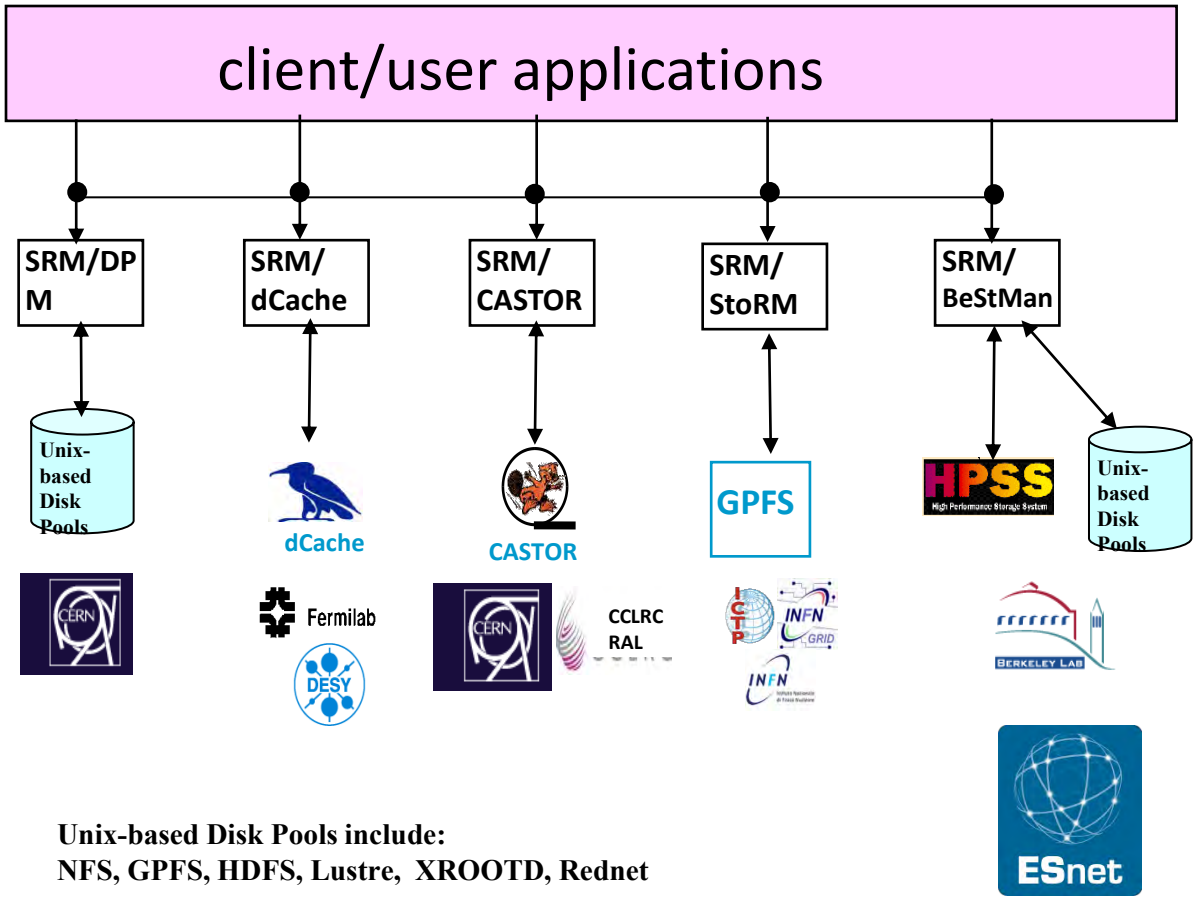

Unix-based Disk Pools include:
NFS, GPFS, HDFS, Lustre, XROOTD, Rednet

Implementation	# of Sites
Bestman/NFS	24
dCache	15
Bestman/HDFS	4
Bestman/XROOTD	4
Bestman/GPFS	2
Bestman/Lustre	2
Bestman/Reddnet	1

Uniform data movement between very different configurations and disk cache sizes.

Data transferred from/to sites in Europe, North America, and Asia.

51 Institutions use inter-operating SRM implementations. Clients can access different mass storage and file systems through a uniform SRM interface.

FastBit - Efficient Search Technology for Data Driven Science

- **Problem**
  - Quickly find records satisfying a set of user-specified conditions in a large, complex data set
  - Example: High-energy physics data –find a few thousand events based on conditions on energy level and number of particles in billions of collision events, with hundreds of variables
- **Solution**
  - Developed new indexing techniques and a new compression method for the indexes, achieved 10-100 fold speedup compared with existing methods
  - Efficient software implementation: available open source from http://sdm.lbl.gov/fastbit/ (1000s of downloads), received a R&D 100 Award
- **Impact**
  - Laser Wakefield Particle Accelerator data analysis: FastBit acts as an efficient back-end for a visual analytics system, providing information for identifying and tracking particles
  - Combustion data analysis: FastBit identifies ignition kernels based on user specified conditions and tracks evolution of the regions
  - Testimonial "FastBit is at least 10x, in many situations 100x, faster than current commercial database technologies" – Senior Software Engineer, Yahoo! Inc