ESnet
ENERGY SCIENCES NETWORK

The Science DMZ: Recent Developments

Eli Dart, Network Engineer
ESnet Science Engagement
Lawrence Berkeley National Laboratory

WRNP17
Belém, Brazil
May 16, 2017

U.S. DEPARTMENT OF ENERGY
Office of Science

BERKELEY LAB

Overview

- Science DMZ As Platform
- Modern Research Data Portal
- Pacific Research Platform
  - PRP
  - NRP

- Note: This talk assumes you already understand the Science DMZ
  - If you haven't encountered the Science DMZ, several folks in RNP can help you, including Leandro Ciuffo and Alex Moura
  - Or check out the fasterdata knowledgebase:
    - http://fasterdata.es.net/science-dmz/

Science DMZ As A Platform

- Once there are many Science DMZs in your network, more things become possible
- Easy file transfer is good, but what else can we do?
  - Update the architecture of data portals
  - Build services between institutions
  - Interconnect facilities
- Several efforts underway to do these things

Science Data Portals

- Large repositories of scientific data
  - Climate data
  - Sky surveys (astronomy, cosmology)
  - Many others
  - Data search, browsing, access
- Many scientific data portals were designed 15+ years ago
  - Single-web-server design
  - Data browse/search, data access, user awareness all in a single system
  - All the data goes through the portal server
    - In many cases by design
    - E.g. embargo before publication (enforce access control)

Legacy Portal Design

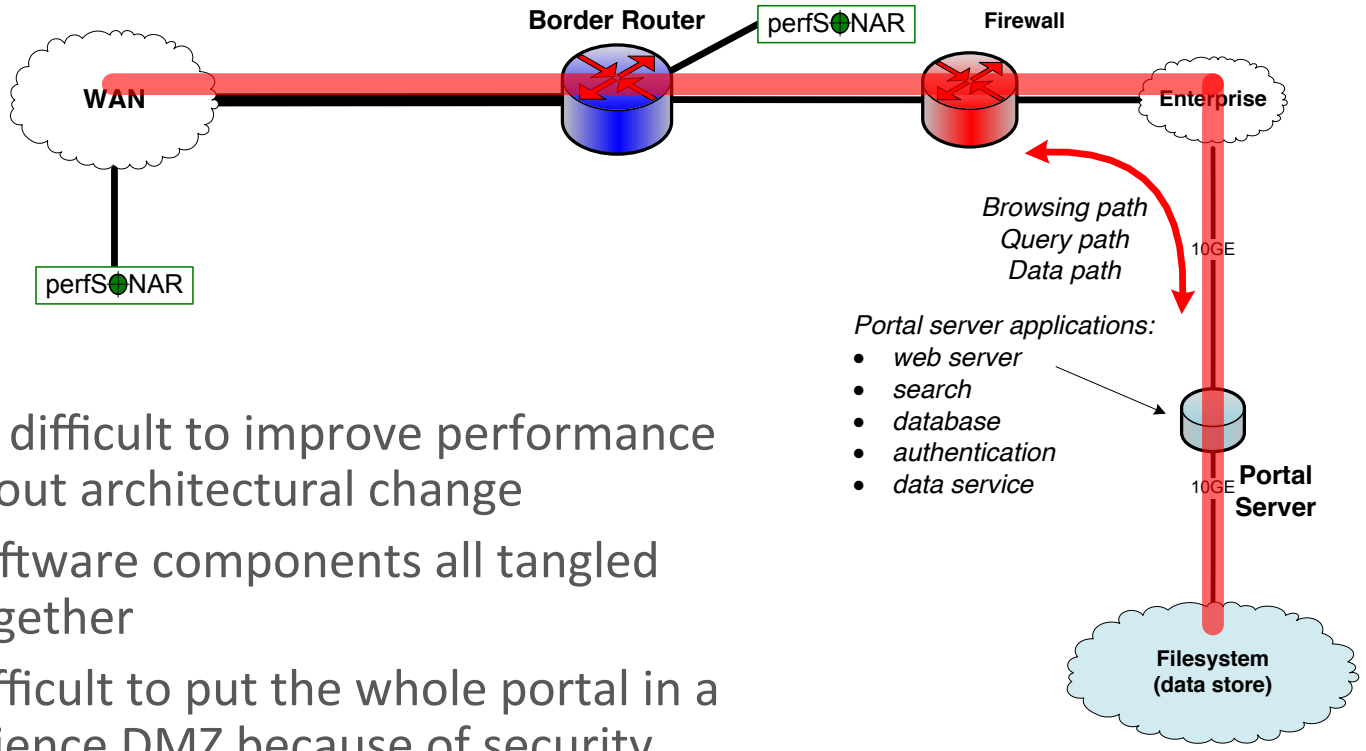

- Very difficult to improve performance without architectural change
  - Software components all tangled together
  - Difficult to put the whole portal in a Science DMZ because of security
  - Even if you could put it in a DMZ, many components aren't scalable
- What does architectural change mean?

Example of Architectural Change – CDN

- Let's look at what Content Delivery Networks did for web applications
- CDNs are a well-deployed design pattern
  - Akamai and friends
  - Entire industry in CDNs
  - Assumed part of today's Internet architecture
- What does a CDN do?
  - Store static content in a separate location from dynamic content
    - Complexity isn't in the static content – it's in the application dynamics
    - Web applications are complex, full-featured, and slow
      - Databases, user awareness, etc.
      - Lots of integrated pieces
    - Data service for static content is simple by comparison
  - Separation of application and data service allows each to be optimized

ESnet

Classical Web Server Model

- Web browser fetches pages from web server
  - All content stored on the web server
  - Web applications run on the web server
    - Web server may call out to local database
    - Fundamentally all processing is local to the web server
  - Web server sends data to client browser over the network
- Perceived client performance changes with network conditions
  - Several problems in the general case
  - Latency increases time to page render
  - Packet loss + latency cause problems for large static objects

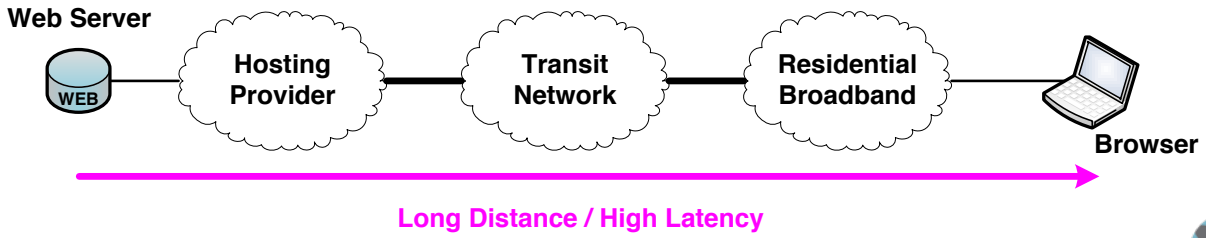

ESnet

Solution: Place Large Static Objects Near Client

- CDN provides static content "close" to client
  - Latency goes down
    - Time to page render goes down
    - Static content performance goes up
  - Load on web server goes down (no need to serve static content)
  - Web server still manages complex behavior
    - Local reasoning / fast changes for application owner
- Significant win for web application performance

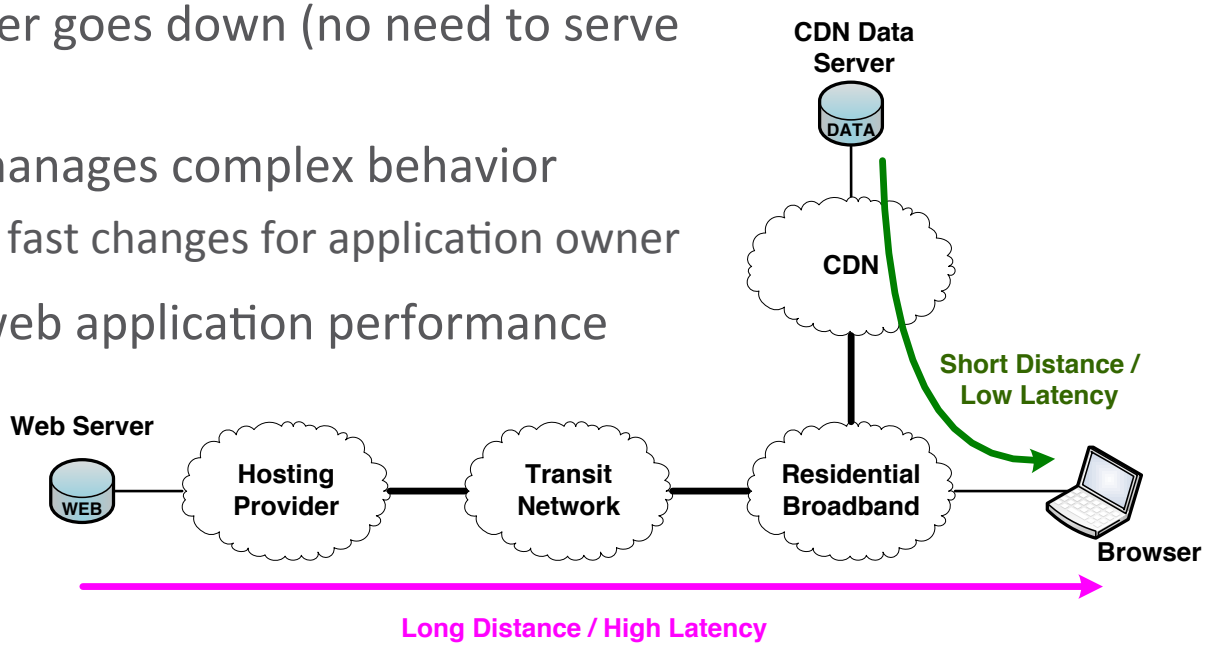

Client Simply Sees Increased Performance

- Client doesn't see the CDN as a separate thing
  - Web content is all still viewed in a browser
    - Browser fetches what the page tells it to fetch
    - Different content comes from different places
    - User doesn't know/care
- CDNs provide an architectural solution to a performance problem
  - Not brute-force
  - Work smarter, not harder

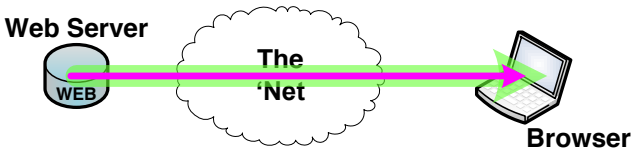

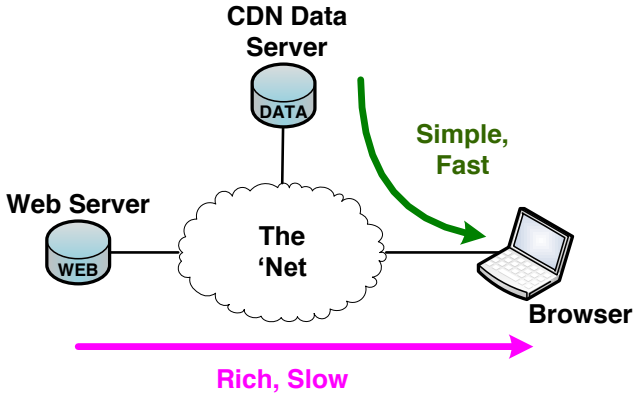

Architectural Examination of Data Portals

- Common data portal functions (most portals have these)
  - Search/query/discovery
  - Data download method for data access
  - GUI for browsing by humans
  - API for machine access – ideally incorporates search/query + download
- Performance pain is primarily in the data handling piece
  - Rapid increase in data scale eclipsed legacy software stack capabilities
  - Portal servers often stuck in enterprise network
- Can we "disassemble" the portal and put the pieces back together better?
  - Use Science DMZ as a platform for the data piece
  - Avoid placing complex software in the Science DMZ

Legacy Portal Design

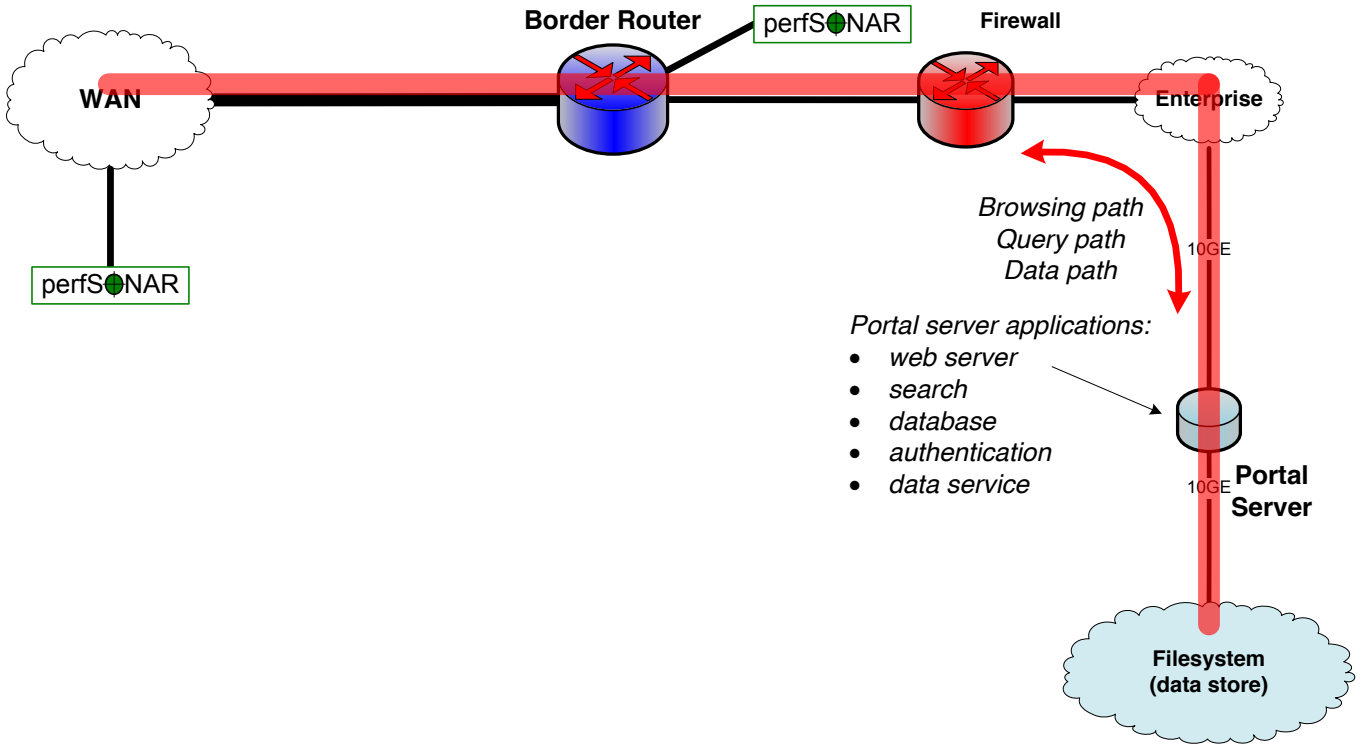

Next-Generation Portal Leverages Science DMZ

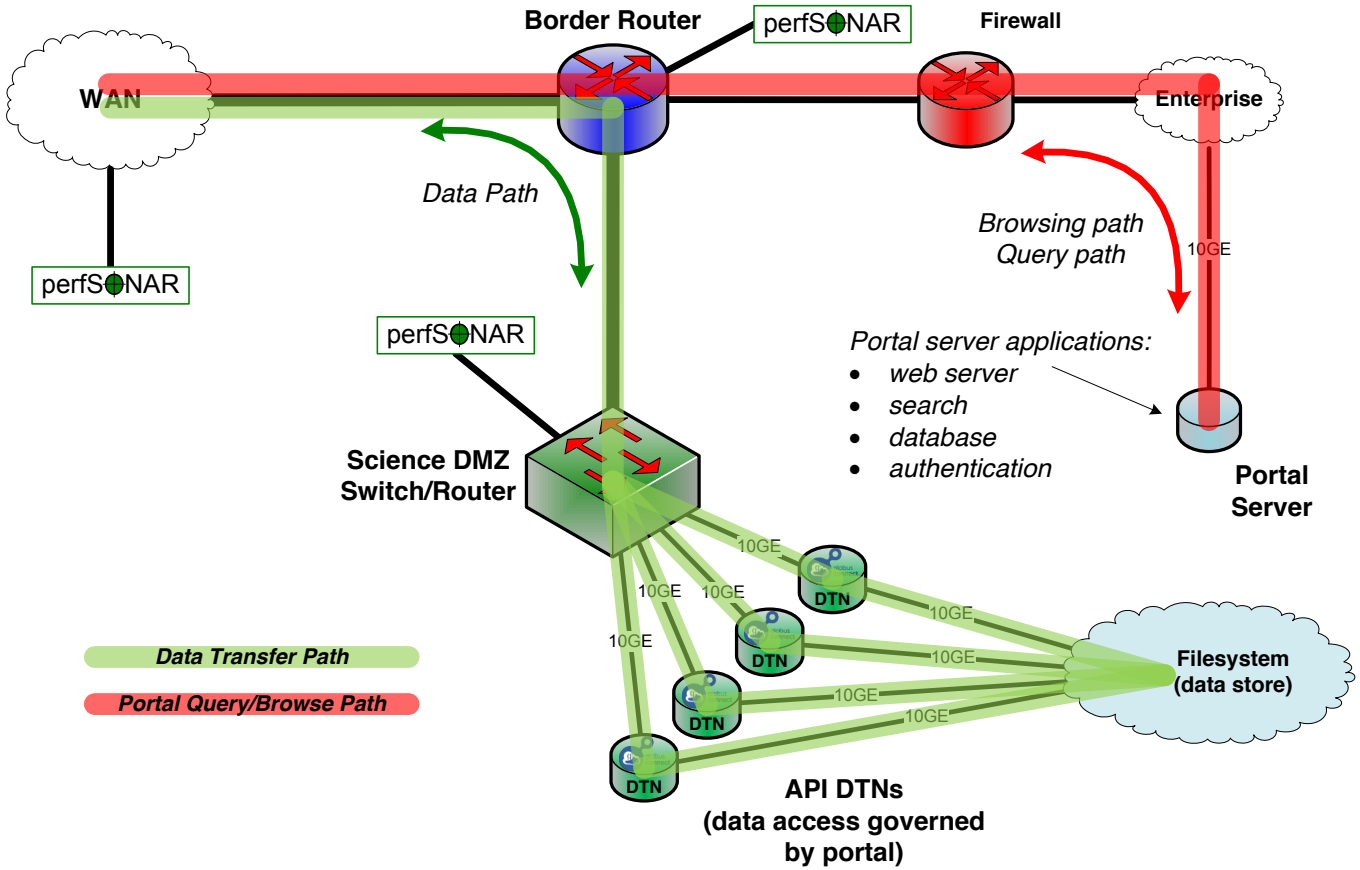

Put The Data On Dedicated Infrastructure

- We have separated the data handling from the portal logic
- Portal is still its normal self, but enhanced
  - Portal GUI, database, search, etc. all function as they did before
  - Query returns pointers to data objects in the Science DMZ
  - Portal is now freed from ties to the data servers (run it on Amazon if you want!)
- Data handling is separate, and scalable
  - High-performance DTNs in the Science DMZ
  - Scale as much as you need to without modifying the portal software
- Outsource data handling to computing centers or campus central storage
  - Computing centers are set up for large-scale data
  - Let them handle the large-scale data, and let the portal do the orchestration of data placement

The Pacific Research Platform Creates a Regional End-to-End Science-Driven "Big Data Freeway System"

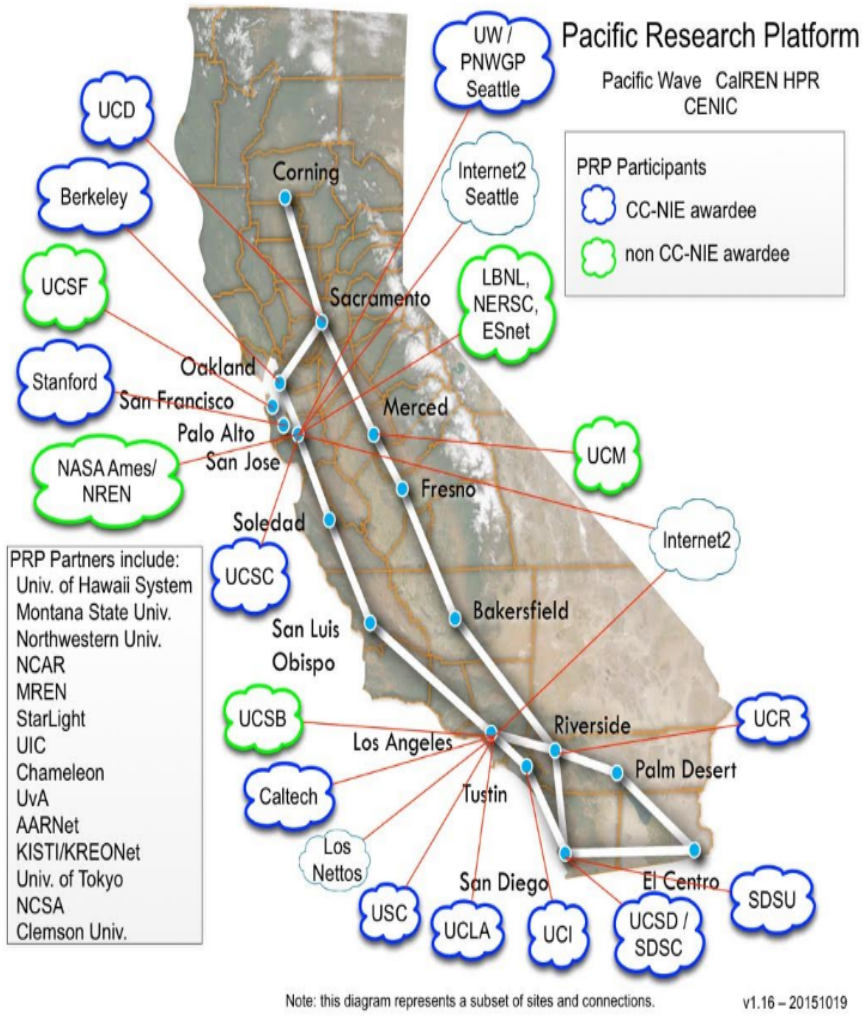

Source:
John Hess, CENIC

NSF CC*DNI Grant
$5M 10/2015-10/2020

- PI: Larry Smarr, UC San Diego Calit2
- Co-PIs:
  - Camille Crittenden, UC Berkeley CITRIS,
  - Tom DeFanti, UC San Diego Calit2,
  - Philip Papadopoulos, UC San Diego SDSC,
  - Frank Wuerthwein, UC San Diego Physics and SDSC

ESnet

PRP Provides Interoperability

- Science DMZs at participating sites ensure interoperability
- PRP engineers work to ensure they interoperate
  - Globus data transfer between DTNs
  - perfSONAR

- Some variation in DTNs
  - Some have FIONA DTNs
    - FIONA == Flash I/O Network Appliance
    - Designed by PRP engineers at UC San Diego
    - https://fasterdata.es.net/science-dmz/DTN/fiona-flash-i-o-network-appliance/
  - Some have DTNs connected to HPC storage
- Key – they all interoperate, removing integration burden from scientists

PRP Science Drivers

- Multiple science areas
  - Astronomy and astrophysics
  - Biomedical applications
  - Life sciences
  - Particle physics
  - Virtual reality and data visualization

- http://prp.ucsd.edu/

ESnet

National Research Platform (NRP)

- Replicate the PRP on a national scale
- Interoperable, high-performance cyberinfrastructure
  - Built to serve domain science
  - Scale up to ~200 institutions
- First workshop to be held this summer
  - Domain science input
  - Policy questions
  - Architecture, scalability
  - Include campus IT, regional networks, national networks, funding agencies, etc. in a common conversation.

Petascale DTN Project

- Another example of building on the Science DMZ
- Supports all data-intensive applications which require large-scale data placement
- Collaboration between HPC facilities
  - ALCF, NCSA, NERSC, OLCF
- Goal: per-Globus-job performance at 1PB/week level
  - 15 gigabits per second
  - With checksums turned on, etc.
  - No special shortcuts, no arcane options
- Reference data set is 4.4TB of astrophysics model output
  - Mix of file sizes
  - Many directories
  - Real data!

ESnet

Petascale DTN Project

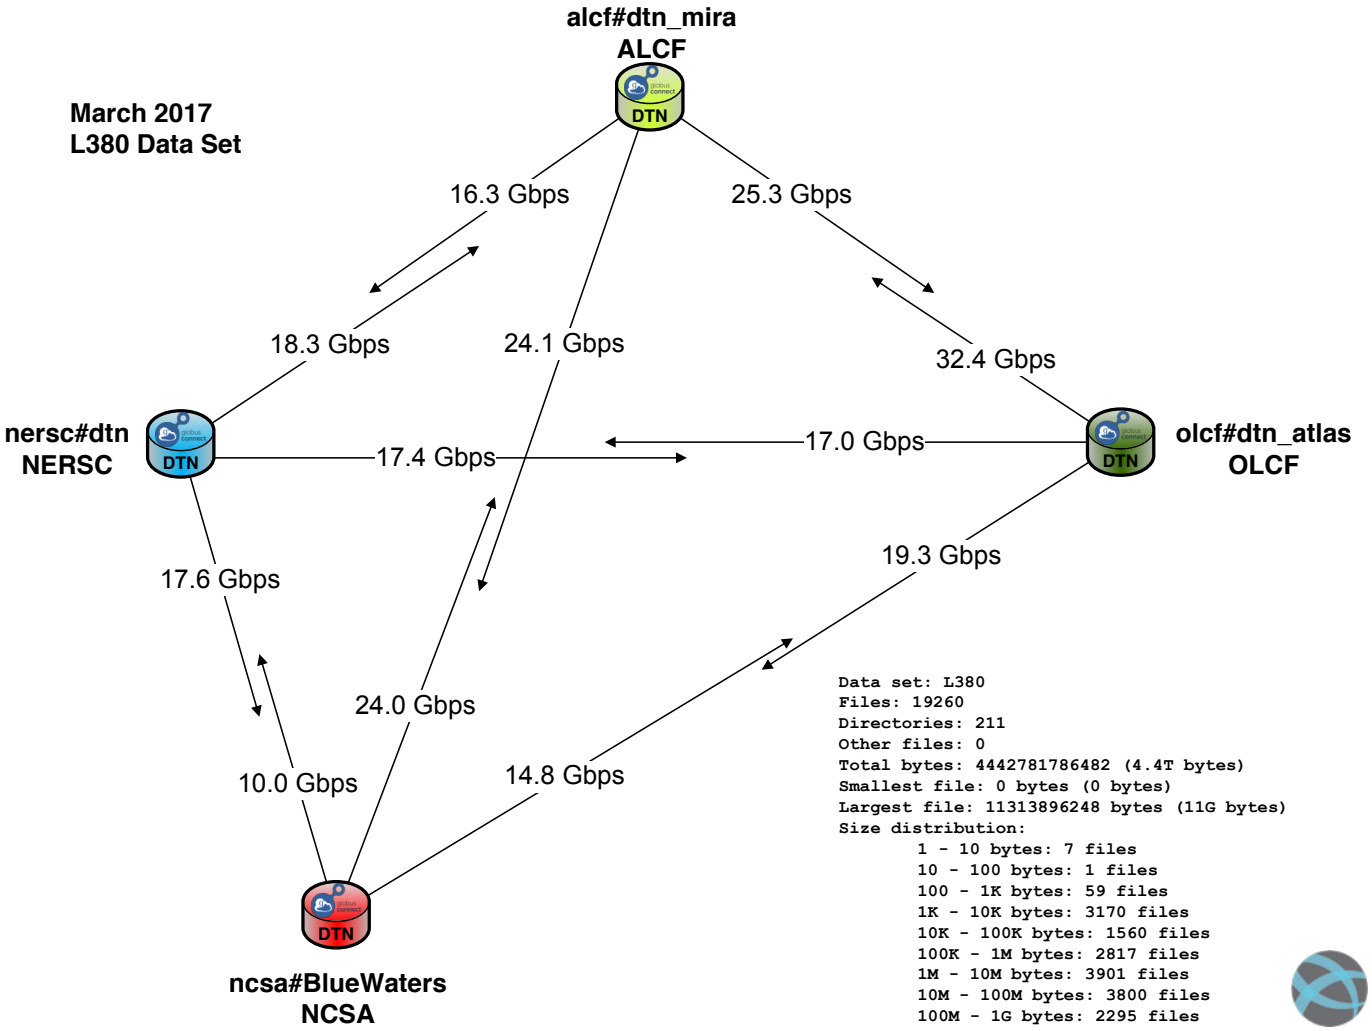


```
Data set: L380
Files: 19260
Directories: 211
Other files: 0
Total bytes: 4442781786482 (4.4T bytes)
Smallest file: 0 bytes (0 bytes)
Largest file: 11313896248 bytes (11G bytes)
Size distribution:
        1 - 10 bytes: 7 files
        10 - 100 bytes: 1 files
        100 - 1K bytes: 59 files
        1K - 10K bytes: 3170 files
        10K - 100K bytes: 1560 files
        100K - 1M bytes: 2817 files
        1M - 10M bytes: 3901 files
        10M - 100M bytes: 3800 files
        100M - 1G bytes: 2295 files
        1G - 10G bytes: 1647 files
        10G - 100G bytes: 3 files
```

ESnet
ENERGY SCIENCES NETWORK

Thanks!

Eli Dart dart@es.net
Energy Sciences Network (ESnet)
Lawrence Berkeley National Laboratory

http://fasterdata.es.net/
http://my.es.net/
http://www.es.net/

Extra Slides

What Is Science Engagement?

- Technology people working with scientists to help solve problems
  - Improve data transfer performance
  - Improve data workflows (e.g. to require less human effort)
  - Improve experiment operations
  - …and more…
- Using experience gained from helping scientists to improve cyberinfrastructure
  - Network design
  - Tool design
  - System design

ESnet

Engagement Is Important: Old Model

- Scientist as integrator
  - Requires scientists to discover new technologies
  - Requires scientists to become expert in new technologies
  - Requires scientists to assemble distinct technologies into an integrated solution that works for them
  - Some scientists do this brilliantly – most do not

ESnet

Engagement Is Important: New Model

- Scientist as collaborator
  - Technologists understand technology
  - Technologists understand enough of the science to see how technology fits
  - Technologists help scientists adopt a useful solution
  - This is much more productive, and requires science engagement

ESnet