

# Experience from HPCs

Doug Benjamin

Soon to be Brookhaven National Lab


**BROOKHAVEN**  
NATIONAL LABORATORY



BROOKHAVEN SCIENCE ASSOCIATES

# The Problem

- Big HPC centers serve wide community of users
- Self Service tools that are easy for users to use are important
- Tools including those for WAN transfers must comply with local Site authentication policy
- Data transfer services must be powerful and scalable



# US HPC WAN Transfer challenge

- Eli Dart – PetaScale DTN Project
- To achieve these data transfer rates – multiple data transfer nodes deployed at each HPC

Petascale DTN Project

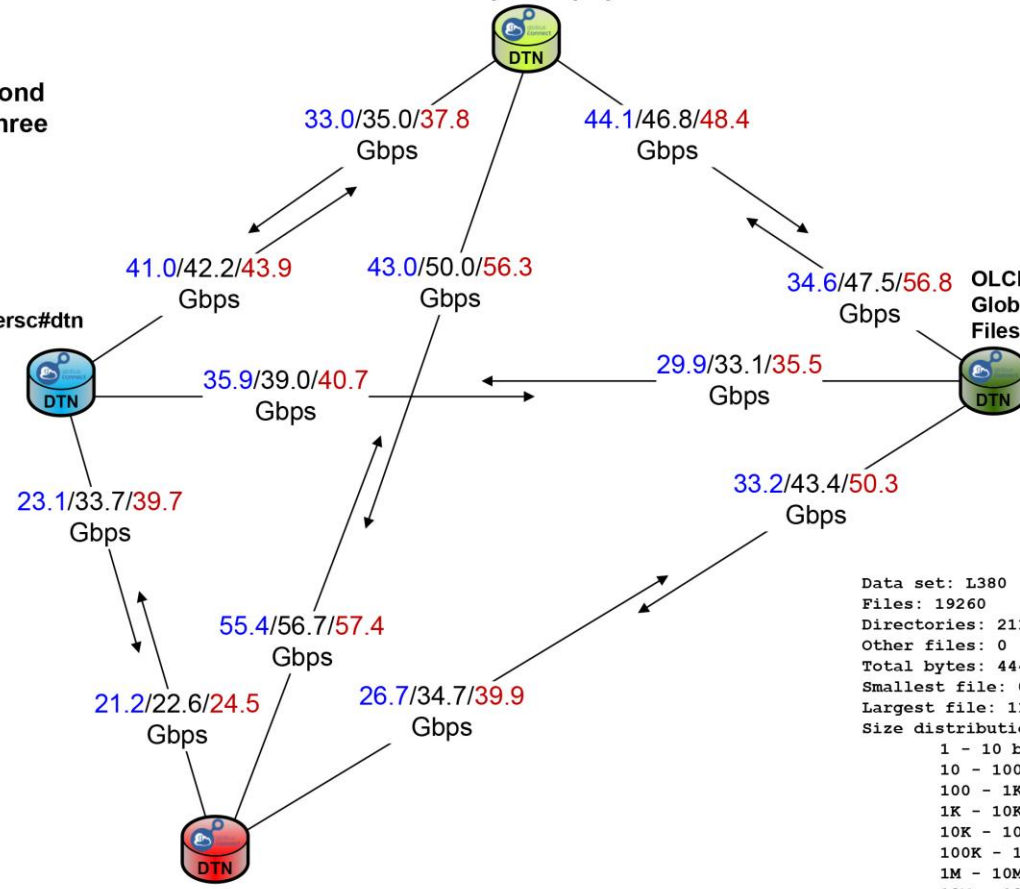
November 2017  
L380 Data Set

Gigabits per second  
(min/avg/max), three transfers

NERSC DTN cluster  
Globus endpoint: nersc#dtn  
Filesystem: /project

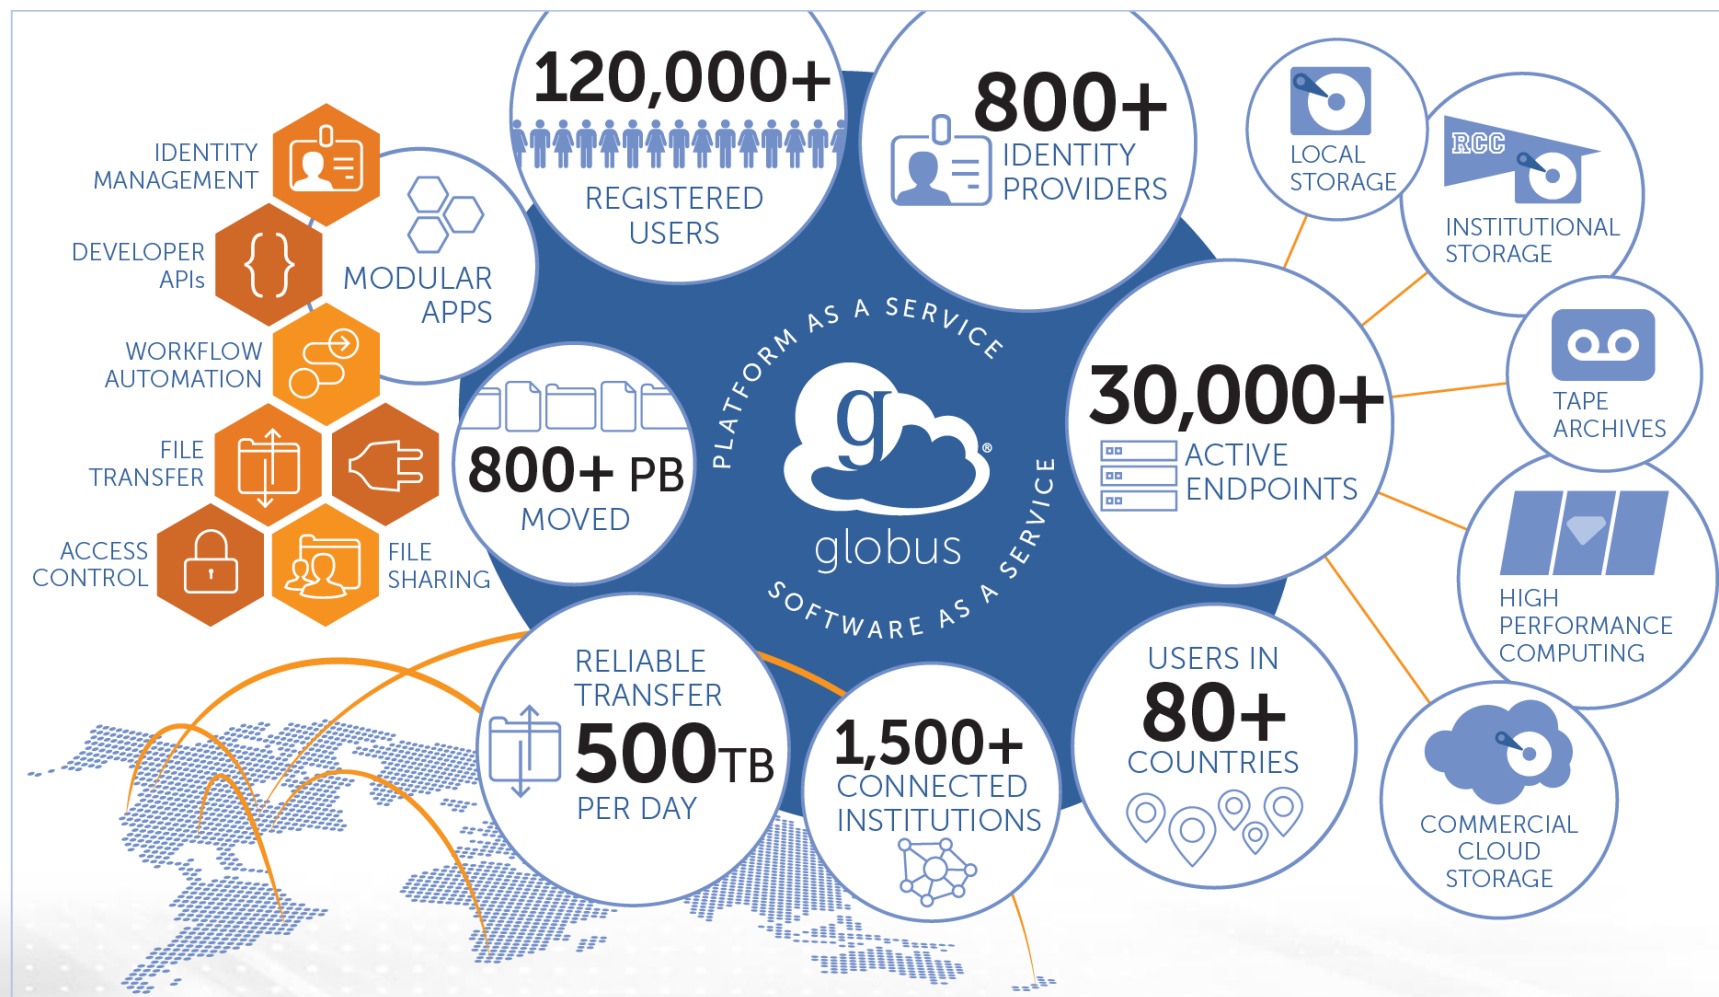
ALCF DTN cluster  
Globus endpoint: alcf#dtn\_mira  
Filesystem: /projects

OLCF DTN cluster  
Globus endpoint: olcf#dtn\_atlas  
Filesystem: atlas2



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

# A Solution for WAN transfers – Globus Service



At a HPC Center in Europe – LRZ (experimentally)

US – All large HPC's

No official discussion from Euro-HPC.

**Not likely to use GLOBUS**



# Easy to use Web interface

- User can manage her own transfers -

The screenshot shows a web interface for managing endpoints. At the top, there is a search bar labeled "Search all endpoints" and a button to "Create a personal endpoint". Below the search bar are several filter tabs: "Recently Used", "In Use", "Shareable By You", "Shared With You", and "Administered By You". A "Filter recently used" dropdown is also present. The main content is a table with the following columns: ENDPOINT, STRICT, STATUS, ROLE, and SHARED. The table lists seven endpoints with their respective icons, names, and details.

ENDPOINT	STRICT	STATUS	ROLE	SHARED
<b>alcf#dtn_theta</b> Managed Public Endpoint		ready		
<b>dcdoor01</b> Managed Public Endpoint		ready		
<b>lcrc#dtn_bebop</b> Managed Public Endpoint		ready		
<b>NERSC DTN</b> Managed Public Endpoint		ready		
<b>SDCC</b> Managed Public Endpoint		inactive		
<b>slac#osg</b> Public Endpoint		ready		
<b>ubuntu01</b> Globus Connect Personal		offline		

Endpoints  
used recently

# Comply with local HPC authentication policies and tools



Argonne Leadership Computing Facility

NERSC – w/ MFA

## MyProxy Client Authorization

Welcome to the OAuth for MyProxy access to your account. If you appr

### Client Information

Name: Globus  
URL: <https://www.globus.org>

Username

Password

MFA Token

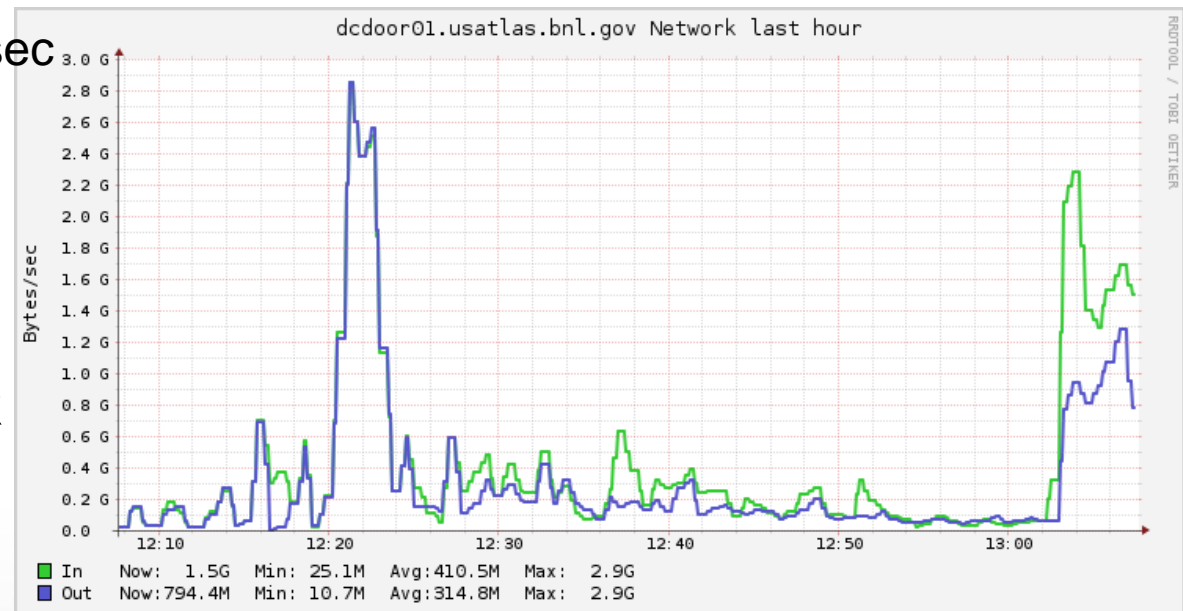
Welcome to the ANL ALCF Authorization Page	
Science Gateway Access	Sign in
The Science Gateway listed below is requesting access to your ANL ALCF account. If you approve, please sign in using your ALCF username along with your CryptoCard PIN+password (for physical token), or CryptoCard password (for mobile token).	
Name: Globus URL: <a href="https://www.globus.org/">https://www.globus.org/</a>	Username <input type="text" value="benjamin"/> Password <input type="password"/> <input type="button" value="Sign In"/> <input type="button" value="Cancel"/>
<small>Please send any questions or comments about this site to <a href="mailto:support@alcf.anl.gov">support@alcf.anl.gov</a>.</small>	

ALCF – w/ Dongle

# Used by ATLAS in US

- Use Globus Python SDK and Harvester to transfer data
- Used extensively in US between BNL and NERSC – Cori HPC and BNL - ALCF Theta and University of Chicago (MWT2) and TACC-Frontera.
- In user for many years
- Not perfect
  - Manual intervention
    - authentication endpoints 1x week
  - **Not integrated w/ Rucio**

3 GBytes/sec



Network traffic gridftp door in 1 hour  
at BNL

# EuroHPC

## (Andrej Filipcic)

- EuroHPC federation plan with hyperconverged network (1Tb/s) to GEANT
- Relying on PRACE MDVPN (similar to LHCONE), although it connects only GEANT, SurfNET, JANET, DFN for now (and HPCs attached there)
- Data transfers not discussed within EuroHPC yet, initially adopting PRACE model (using ssh)
- no technical details or policies yet
  - (primary focus is to commission the machines in 2021)

***European HPC's and US HPC's will have different solutions***



# Summary

- Globus provides easy to use service to transfer data between HPC centers and their users
  - NERSC Users use it to move a lot of data internally w/ NERSC
- Scalable – Complies with local authentication policies
- Manages transfers
- Many Sites are subscribers
- **Used in US extensively**
- **Not w/ Europe-HPC**

The screenshot shows the Globus website interface. At the top is a dark blue navigation bar with the Globus logo (a cloud with a 'g') and the text "globus a uchiicago non-profit service". To the right of the logo are links for "I Want To...", "Pricing", "Resources", "Support", "About", and a "Log In" button. Below the navigation bar is a white section with the heading "Transfer your data." in blue. Underneath is a diagram showing a cloud with the Globus logo at the top, connected by dashed lines to two blue circles labeled "A" and "B". A thick orange arrow points from "A" to "B", with the text "fast secure transfer reliable" written across it. Below the diagram is a paragraph of text: "Gigabytes, terabytes, petabytes—research data is large and distributed. Globus lets you efficiently, securely, and reliably transfer data directly between systems separated by an office wall or an ocean. Focus on your research and offload your data transfer headaches to Globus." To the right of this text are three blue links: "LEARN MORE", "TRANSFER DATA NOW", and "CONNECT YOUR SYSTEM TO GLOBUS". At the bottom of the screenshot is a large digital counter displaying "1,122,721,450,548 MB TRANSFERRED".