


NUG Teleconference November 2013



Richard Gerber
NERSC User Services

November 7, 2013

Connection Info



Connection Info

Topic: NUG Web Conference

Date and Time:

Thursday, Nov. 7, 2013 11:00 am, Pacific Standard Time

Event number: 663 554 441

Event password: edison

<https://nersc-training.webex.com/>

Teleconference information

1-866-740-1260

PIN: 4866820

- **Edison Update**
- **NUG 2014 Annual Usage Group Meeting**
- **2014 NISE Call Coming**
- **NERSC Achievement Awards**
- **Queue Committee**
- **NERSC “Quarterly” Outage**
- **Using NERSC’s Data Transfer Nodes**
- **Additional items and Q&A**

- **Edison continues to deliver hours to science**
 - 257 million hours in 2013 (514 million Hopper equivalent)
- **Still undergoing system evaluation and testing**
 - DARPA mission partners (11 million hours)
 - NERSC
 - Cray
- **Outstanding issue**
 - node failure rate higher than expected

NUG 2013 Annual Meeting



- **Feb. 4-6 in Berkeley/Oakland**
- **NUG committee is working with NERSC and Berkeley Lab**
 - Lead by Frank Tsung
- **Kickoff of NERSC 40th Anniversary celebration**
- <http://www.nersc.gov/users/NUG/annual-meetings/nug-2014/draft-agenda/>
- **Proposed Topics**
 - Where has NERSC been and where is it going?
 - Achievement awards presentations & talks
 - Science keynote(s)
 - HPC technology outlook
 - How will we program future systems?
 - What do scientists want?
 - Special theme topics

- Call for proposals will be issued soon.
- Proposals will be submitted through the ERCAP interface.
- 100 million NERSC MPP hours will be allocated for 2014.
- Scope
 - HPC and data analysis: Projects that require extreme scale computing to analyze big data sets.
 - Highly collaborative HPC: Projects that join community teams, knowledge bases or gateways with HPC computing and data capabilities.
 - Extreme networking: Large-scale computing projects that leverage NERSC's 100GB border network connectivity.

- **Committee of users from NUGEX**
 - Determine nomination process
 - Make evaluations and recommendations with input from domain experts
 - Being lead by Stephane Ethier and Cameron Geddes
 - Still need representatives from “Big Data,” biosciences
- **Call for nominations will be issued next week**


- **NUG Committee is forming to address HTC and general queue and throughput needs**
- **NERSC will facilitate, but we'd like it to be driven by NUG**
- **I'll set up a kickoff teleconference next week, but then it's up to you all**
- **If you would like to volunteer to be on the committee, please do so!**
 - Committee will be chosen to represent a cross-section of NERSC users
 - Email to ragerber@nersc.gov

NERSC Quarterly Maintenance



- **Date: 2013/11/12**
- **Systems affected**
 - **Mendel** (Matgen, Planck, Genepool Phase2) –
 - **DOWN FOR WEEK to upgrade network reliability and performance.**
 - 11/12/13 19:00-11/19/13 19:00 PST
 - JGI – Genepool (and other JGI related systems)
 - 11/12/13 7:00-19:00 PST
 - Rolling outages on network, LDAP, Science Gateways
 - Expected to only last a few minutes for a given service
 - Be aware of this if your jobs depend on other services at NERSC

Data Transfer Nodes at NERSC



Shreyas Cholia
Outreach, Software and Programming Group

November 7, 2013

What are the DTNs?



- **The Data Transfer Nodes (DTN) are servers dedicated to data transfer at NERSC.**
 - Nodes – `dtn[01-04].nersc.gov`
- **DTNs have access to most of the NERSC file systems, and are tuned to transfer data efficiently.**
- **The Data Transfer Nodes are tuned for transferring large volumes of data between NERSC and other major facilities (ORNL, ANL etc.)**
- **Can also be used to move data between NERSC file systems and HPSS**

In short ...



- **Use the DTNs if you want to move large volumes of data in and out of NERSC (or between NERSC systems)**

Login Access



- **All NERSC users have login access**
- **NERSC Users (non-JGI):**
 - `ssh dtn01.nersc.gov` (or `dtn02`)
- **JGI Users:**
 - `ssh dtn03.nersc.gov` (or `dtn04`)
- **Familiar module environment**
 - `module avail`

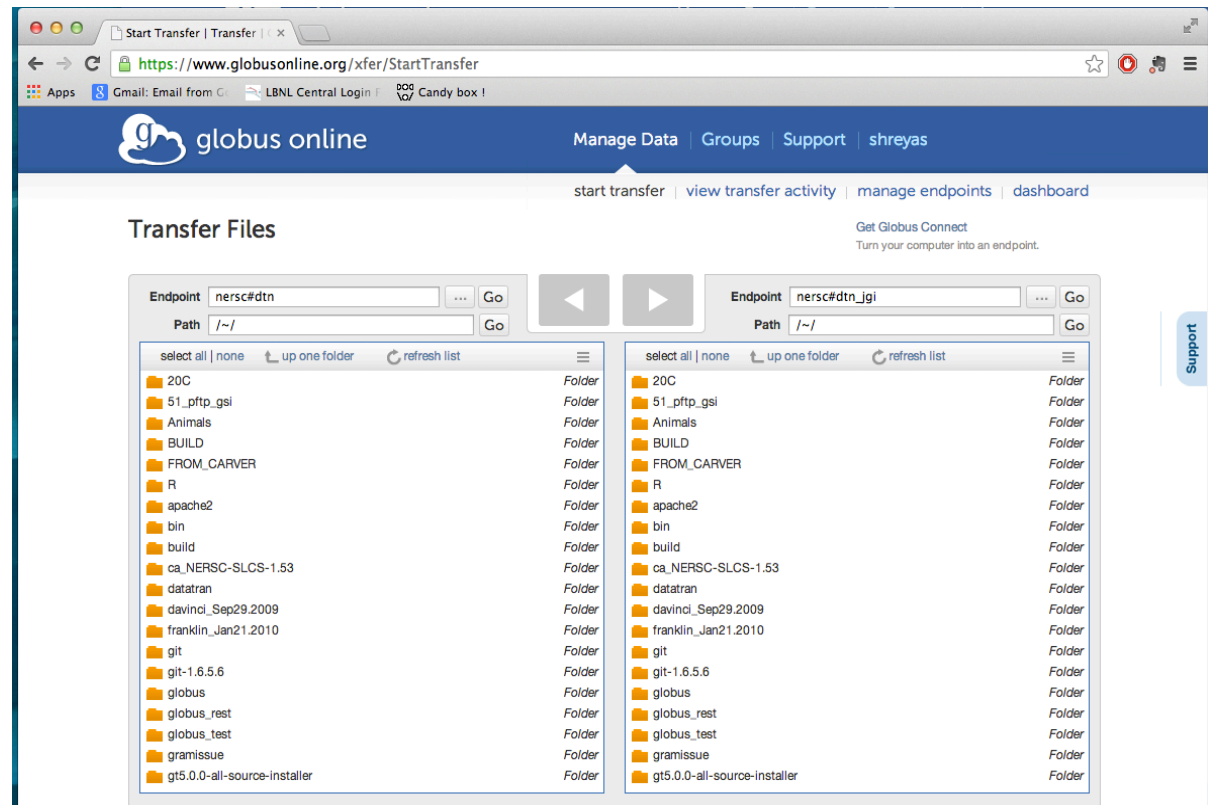
- All Global Filesystems
- /global/homes
- /global/scratch2
- /global/project
- /global/projectb
- /global/dna
- /global/seqfs
- /global/common

Transfer Tools



- **Globus Online**
- **scp**
- **bbcp**
- **hsi/htar (for HPSS)**
- **GridFTP**

- Managed 3rd party Transfers – <http://globusonline.org>
- CLI also available – `ssh cli.globusonline.org`



- <http://globusonline.org>
- Select `ner-sc#dtn` endpoint in the transfer window (autocomplete might take a while)
- Select remote endpoint
- Use point and click interface to submit transfer
- Transfer will happen in the background (including retries on failure).
- You can check back on status later
- Globus Connect allows you to have a private endpoint on your laptop

- **These are command line tools – use these directly from the DTN nodes or from a remote node**
 - scp
 - bbcp
 - globus-url-copy

- **Detailed instructions, syntax etc.**
 - <http://www.nersc.gov/users/data-and-file-systems/transferring-data/>

- **To backup or archive your data use HSI and HTAR**
 - <https://www.nersc.gov/users/data-and-file-systems/hpss/storing-and-retrieving-data/>
- **Login to DTN node and run hsi/htar**
 - Preferred to doing this from interactive nodes for long running transfers
- **Can also to Globus Online transfer**
 - `nersc#dtn <-> nersc#hpss`

Why DTNs?



- **Science DMZ**
- **No firewall restrictions**
- **Tuned for WAN transfers**
 - Fast network (ESnet), optimized configuration
- **You don't get booted for long running transfers**
- **Fixed endpoint (in case you need to tune firewalls on the other end)**
- **Dedicated support for data transfer**

- **Use GO for large automated or monitored transfers**
- **scp should be fine for smaller transfers (<100MB)**
- **Expected Performance**
 - This varies from one site to the other.
 - 100-300 MB/sec on a 100G connected site (ANL)
 - 20-100 MB/sec on a 10G connected site
 - Disk may be the bottleneck at high speeds
 - Generally: 1TB per 8 hours is reasonable
- **Bottom Line: If you don't think you are getting the performance you expect, let us know**
 - `consult@nersc.gov`

- **Performance is most often gated by the remote endpoint**
 - they often are not tuned for WAN transfers
 - often have a 1Gb/sec link.
 - These will lower performance < 100 MB/sec.
- **File system contention may be an issue**
 - Isn't much you can do except try the transfer at a different time or on a different FS.
 - You may need to consider the cost of lost time vs. transferring at a lower rate
- **Don't use global/homes! Instead use \$PROJECT or \$GSCRATCH**

For more information



- **General DTN info**

- <https://www.nersc.gov/systems/data-transfer-nodes/>

- **Data transfer info**

- <http://www.nersc.gov/users/data-and-file-systems/transferring-data/>

Other Issues or Questions?



- **NUG Teleconference**

- Next scheduled: Thu. Dec. 7, 2013
- Send suggested topics and comments to ragerber@lbl.gov

- **NERSC HPC Brown Bag**

- Today @ noon PST
- Same remote connection info
- **XRootD In Perspective, Andrew Hanushevsky, SLAC**

This talk will introduce XRootD and show many of the features that have made it a popular data access solution in HEP and Astrophysics (e.g. ALICE, ATLAS, CMS, and Fermi, among others). The talk will also show how it can be used in I/O restricted exa-scale HPC environments to provide seamless data access.



National Energy Research Scientific Computing Center