

September 26,
2014

XSEDE Campus Bridging Tools

Jim Ferguson

National Institute for Computational Sciences

jwf@utk.edu

XSEDE

Extreme Science and Engineering
Discovery Environment



Quick Advertisement: Student Programs

- Research Experience
 - 12-18 students per summer, 10 weeks
 - Work with researcher / mentors on real projects
 - Includes travel for training and to XSEDE15.
 - \$5K stipend
- XSEDE Scholars
 - Annual program, 30-40 students per year
 - Targets under-represented groups from all institutions
 - Specific training provided on HPC, research
 - Includes XSEDE15 travel



For Student, Educator, Workshop news...

- XSEDE website, Education & Outreach
 - Now under 'Education & Outreach Blog' but soon to reorganized under 'Students' section
- Facebook: Computational Science Education News
- hpcuniversity.org website




What is Campus Bridging?

- “Bridging” the gap between local researcher cyberinfrastructure, campus CI, and national CI resources
- It is not just one thing
 - Hardware
 - Software
 - Support
 - Outreach

Adequacy of research CI

5



Responses to asking if researchers had sufficient access to cyberinfrastructure resources – survey sent to 5,000 researchers selected randomly from 34,623 researchers funded by NSF as Principal Investigators 2005-2009; results based on 1,028 responses

Stewart, C.A., D.S. Katz, D.L. Hart, D. Lantrip, D.S. McCaulay and R.L. Moore. Technical Report: Survey of cyberinfrastructure needs and interests of NSF-funded principal investigators. 2011. hdl.handle.net/2022/9917



XSEDE

More on Campus Bridging

- The goal of campus bridging is virtual proximity ...
- The biggest problems:
 - Not enough CI resources available to most researchers
 - When you go from your campus to the national cyberinfrastructure it can feel like you are falling off a cliff! That's why you need bridging....
- Campus bridging is a major priority within XSEDE



XSEDE Campus Bridging

- Even for those researchers with enough computing time, there has traditionally been little support once the allocation process is complete
 - XSEDE has prioritized training and documentation, as well as a number of initiatives designed to make the transition appear seamless to the user:



Campus Bridging services today

- UNICORE – Grid system developed in Europe, supported and continually evolving, adopted by XSEDE for optional deployment on level 1 and 2 Service Providers. Nearly universal.
- XSEDE Compatible Basic Cluster: Rocks Rolls and YUM Repo
- Globus Online – File transfer utility, approved for use on XSEDE resources. Globus Connect Server also through engineering process.
- Genesis II – global file system (GFS) and job submission service under evaluation & development

UNICORE

- Production CI suite
- Installation instructions at <https://portal.xsede.org/software/unicore>
- Server and Client software available, web interface under development
- Under continued development at Forschungszentrum Jülich in Germany, one developer on XSEDE team.

Rocks Rolls + YUM Repository

- Tools to allow cluster administrators to create a XCBC: the “XSEDE-compatible basic cluster software stack”
- Make your cluster more XSEDE-like
 - Eases later transition to larger CI
 - Saves investment on development and updates
 - Allows researchers to benefit from XSEDE-wide training
 - Teaches students interested in HPC to operate within a context similar to XSEDE



Rocks Rolls + YUM Repository 2

- Maintained and updated by XSEDE Campus Bridging
- New cluster? Rocks ISOs can automate much of the work of setting up a basic XSEDE-like cluster

Rocks Rolls + YUM Repository 3

- Existing cluster? XSEDE provides documentation to configure your cluster such that you can use these resources
- Install YUM Repo locally, manage XSEDE packages as if they were part of the OS, just like you would with the base or EPEL repositories



GlobusOnline

- Simple, Dropbox-like interface for moving files back and forth
- Can be configured to connect between multiple computers and larger resources to transfer data
 - Screenshot courtesy of http://genome.jgi.doe.gov/pages/portal_apps.jsf

The screenshot displays the Globus Online web interface. At the top, the logo and navigation links (Manage Data, Groups, Support, jdoe) are visible. Below the navigation, there are links for 'start transfer', 'view transfer activity', 'manage endpoints', and 'dashboard'. The main section is titled 'Transfer Files' and includes a 'Get Globus Connect' button with the text 'Turn your computer into an endpoint.' Below this, two endpoint selection panels are shown. The left panel has 'Endpoint' set to 'jgi#portal' and 'Path' set to '/~/By_Proposal_Name/~/Acidic_Mine_Environ'. The right panel has 'Endpoint' set to 'nersc#dtn' and 'Path' set to '/~/data/'. Both panels have 'Go' buttons. Below the endpoint panels are two file browser windows. The left window shows a list of folders and files, including 'Acidic_Mine_Environmental_sample__Acid_mine_drainage_[Acimindrainage]', 'Acidic_Mine_Environmental_sample__Community_UBA_Waterfall_[ComUBA]', 'Acidic_Mine_Environmental_sample__Iron_Mountain_AMD_Site_1_[IroMouAM]', 'Acidic_Mine_Environmental_sample__Iron_Mountain_AMD_Site_2_[IroMouAM]', 'Acidic_Mine_Environmental_sample__Ultra_back_A_BS_[UltrabackABS]', 'Acidic_Mine_Environmental_sample__Ultra_back_A_UBA_[UltrabackAUBA]', and 'Acidic_Mine_Environmental_sample__Ultra_back_C_level_1_20_m_back_[Ultraback]'. The right window shows a single folder named 'Acidic_Mine'. At the bottom of the interface, there is a 'Label This Transfer' field with a 'more options' link and a note: 'This will be displayed in your transfer activity.'

Globus Online 2

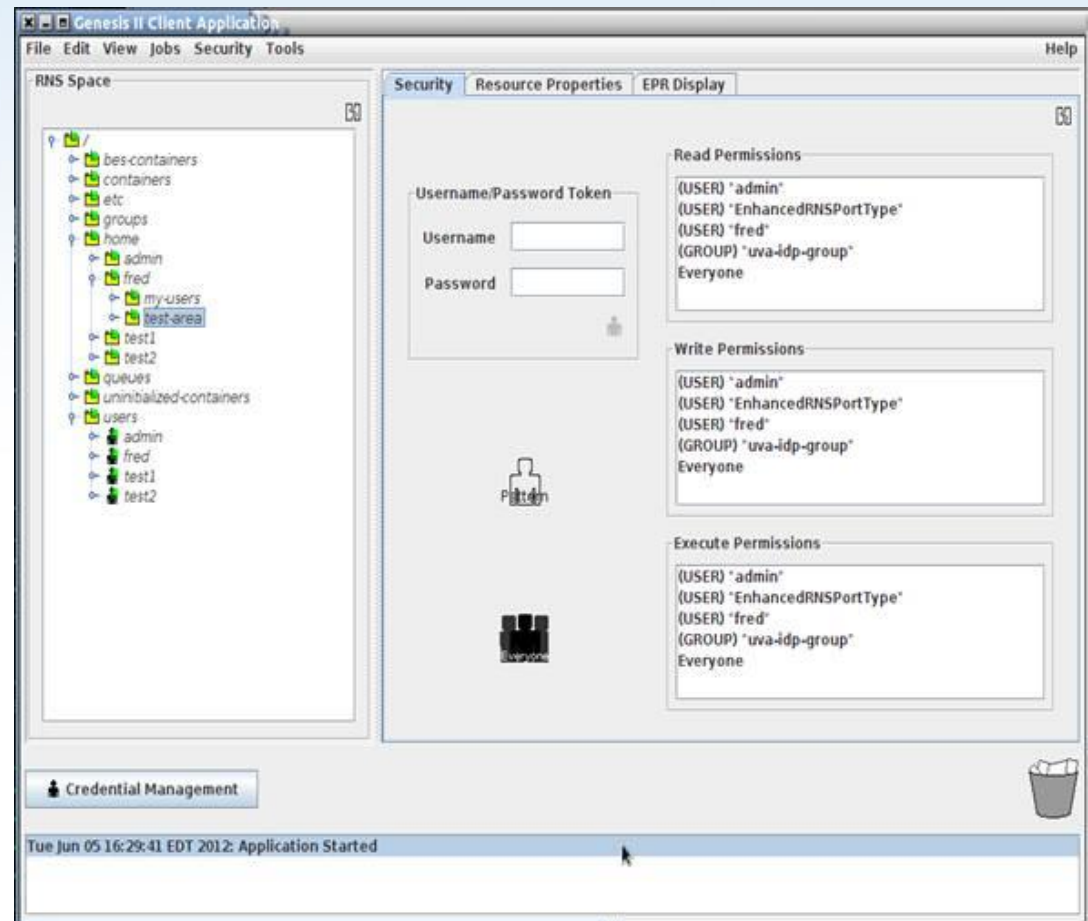
- Set a transfer and start it, walk away, and get a notification when the transfer is complete
- During Q3 2013, GlobusOnline allowed researchers to move over 630TB of data onto the XSEDE network
 - That's just ONE direction!

Genesis II

- Two components
 - Unicore 6 Basic Execution Services (BES)
 - Allows for jobs to be submitted as files into a directory
 - Directory is a jobs queue and BES executes job files as jobs in a queue
 - Global Federated File System (GFFS)
 - Standardizes user interface across compute resources
 - GUI component allows drag and drop uploading, user and group management, and POSIX-style permissions management independent of the command line

GFFS

Screenshot
courtesy
<http://genesis2.virginia.edu>



GFFS 2

- Allows users to export a directory into a globally federated file system that is visible from any system which is running a GenesisII client
- Provides a fuse-based filesystem, mountable in userspace with hooks to job management

Summary

- Campus bridging aims to create virtual proximity between researchers and resources
- Streamlining and easing the analysis part of research
- Ways we are working to bridge research computing:
 - XSEDE-compatible basic cluster stack
 - Full-time dedicated support and consulting
 - GlobusOnline
 - Unicore 6
 - GFFS

For more information...

- <https://www.xsede.org/campus-bridging>
 - <http://cb-repo.iu.xsede.org/xsederrepo/>
 - <https://www.xsede.org/globus-online>
 - <http://genesis2.virginia.edu>
 - campusbridging@xsede.org
 - +1 (812) 318-2872 (M-F, 9a-5p Eastern)
-
- Any questions?



Our reach will forever
exceed our grasp, but,
in stretching our horizon,
we forever improve our world.

XSEDE

Extreme Science and Engineering
Discovery Environment