

BLUE WATERS

SUSTAINED PETASCALE COMPUTING

2014 Blue Waters Update

Bill Kramer

Blue Waters Director



GREAT LAKES CONSORTIUM
FOR PETASCALE COMPUTATION

CRAY®

Announcements

- Today, Ed Seidel has invited the PIs to lunch in the Alma Mater Room.
 - The PI's have blue tickets in the back of their badges.
- We will take a group photo of all attendees at the first break.
- **#BWsymp2014** for another chance



Joint Dinner at Memorial Stadium – Tonight

Joint with the Private Sector Program Workshop Attendees



SETAC

NSF PRAC

- Paul Woodward, Physics and Astrophysics, University of Minnesota
- Tom Cheatham, Chemistry, University of Utah
- Patrick Reed, Civil and Environmental Engineering – Systems Optimization, Cornell
- Klaus Schulten, Physics and Molecular Dynamic, University of Illinois Urbana-Champaign
- David Ceperley, Physics and Material Science, University of Illinois Urbana-Champaign
- Tiziana Di Matteo, Physics and Cosmology, Carnegie Mellon University
- Dave Randall, Atmospheric Sciences and Climate Colorado State University

GLCPC Chair

- Joe Paris, Academic & Research Technologies in Information Technology, Northwestern University (Chair for 2013/2014, followed by Jorge Vinals, Structural Mechanics and Biophysics, University of Minnesota, Chair for 2014/2015)

University of Illinois at Urbana-Champaign Allocation Chair


- Athol Kembal, Atmospheric Sciences, University of Illinois at Urbana-Champaign

Industry

- Rick Authur, General Electric Global Research, Computer and Software Engineering

Blue Waters Fellows

- 6 Awards (so far)
 - Substantial Stipend + Blue Waters allocations
 - 10 other very deserving nominees are being offered Blue Waters allocations
- Kenza Arraki, New Mexico State University
- Jon Calhoun, University of Illinois at Urbana-Champaign
- Sara Kokkila, Stanford University,
- Edwin Mathews, University of Notre Dame
- Ariana Minot, Harvard University
- Derek Vigil-Fowler, University of California, Berkeley



Blue Waters Usage 2/11/14 – Largest 10 Jobs-Torus View

Placement of 10 largest running jobs on the Gemini torus.

Tue 11.02.2014 at 09:39:51 AM CST

JOBID	USERID
-------	--------

574728	jtao
--------	------

576980	redwards
--------	----------

576982	redwards
--------	----------

576985	redwards
--------	----------

576987	redwards
--------	----------


589495	yanxinl
--------	---------

590655	yanxinl
--------	---------

588655	wdaughto
--------	----------

592154	leeping
--------	---------

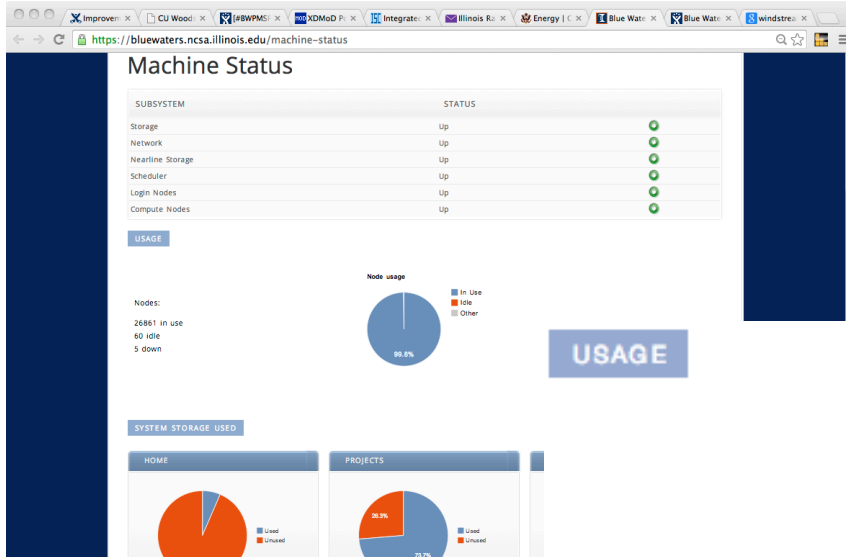
590536	fdm
--------	-----




Each dot is a Gemini router and represents 64 AMD integer cores.

Jobs are 4,096, 2,048 and 1,024 nodes

Blue Waters Usage 2/11/14 – 99.82% of nodes in use




Nodes:
26861 in use
60 idle
5 down



1,024 nodes

Petascale plasma physics simulations using PIC codes (jobs:5) PI: Warren Mori, University of California, Los Angeles	XE	196608	722,907.59
Hierarchical molecular dynamics sampling for assessing pathways and free energies of RNA catalysis, ligand binding, and conformational change (jobs:99) PI: Thomas Cheatham, University of Utah	XE	31680	510,417.29
Hierarchical molecular dynamics sampling for assessing pathways and free energies of RNA catalysis, ligand binding, and conformational change (jobs:45) PI: Thomas Cheatham, University of Utah	XK	31280	178,769.30
Petascale Multiscale Simulations of Biomolecular Systems (jobs:2) PI: Gregory Voth, University of Chicago	XK	10240	

CURRENT RUNNING JOBS BY SCIENCE AREA



Q1 14 Accounts and Projects

Active S&E Projects

- NSF PRACs - 28
 - 5 have completed
- NSF Startup – 6
- GLCPC - 9
 - 8 have completed
- Illinois – 36
- Blue Waters Professors – 11
- Education – 3
 - 6 Fellows
 - 3 have completed
- Innovation and Exploration – 7
- Private Sector Program - 4

Types of S&E Partners

- Faculty: 92
- Researcher: 117
- Postdoctorate: 100
- Graduate Student: 249
- Undergraduate Student: 17
- Total: 575

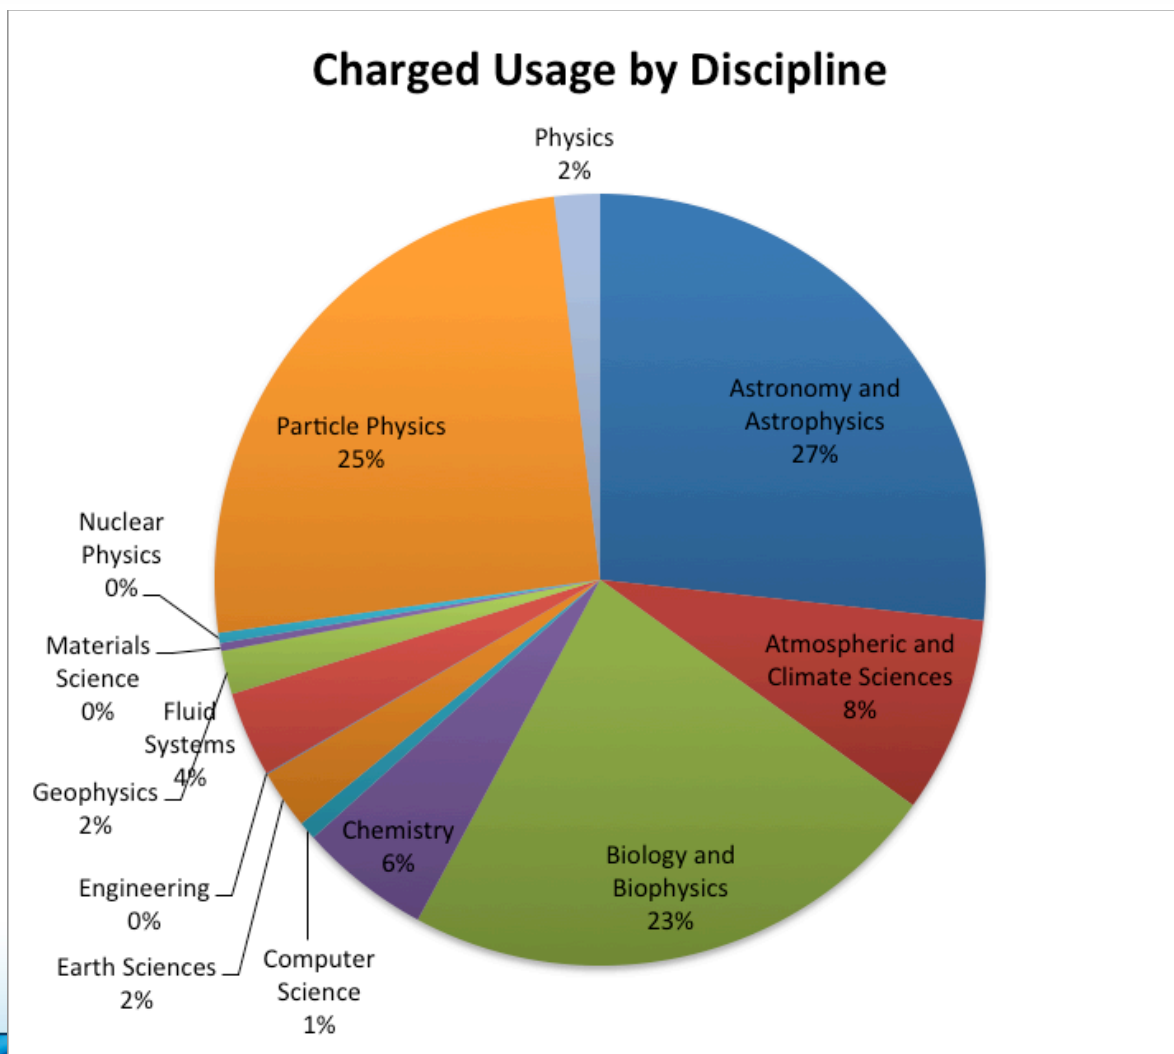
S&E Researchers from 56 Universities

Plus National Labs, Research Centers, industrial organizations


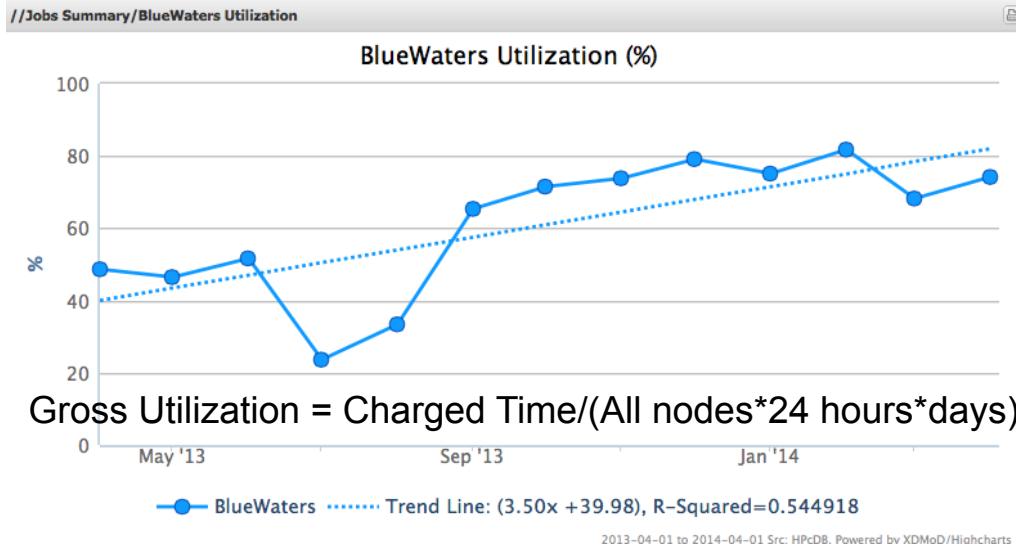
- Brigham Young University
- California Institute of Technology
- Carnegie Mellon University
- Central Michigan University
- College of William and Mary
- Colorado State University
- Columbia University in the City of New York
- Cornell University
- Georgia Institute of Technology
- Indiana University
- Louisiana State University
- Michigan State University
- North Carolina State University at Raleigh
- Pennsylvania State University
- Princeton University
- Purdue University
- Rochester Institute of Technology
- Rutgers, the State University of New Jersey
- SUNY at Stony Brook
- San Diego State University
- Southern Methodist University
- Stanford University
- University of Alabama, Huntsville
- University of Arizona
- University of California, Berkeley
- University of California, Davis
- University of California, Irvine
- University of California, Los Angeles
- University of California, San Diego
- University of California, San Francisco
- University of California, Santa Barbara
- University of California, Santa Cruz
- University of Chicago
- University of Copenhagen
- University of Florida
- University of Granada
- University of Houston-Clear Lake
- University of Illinois at Chicago
- University of Illinois at Urbana-Champaign
- University of L'Aquila
- University of Maryland
- University of Maryland, College Park
- University of Memphis
- University of Michigan
- University of Minnesota
- University of Nevada-Las Vegas
- University of New Hampshire
- University of Notre Dame
- University of Oslo
- University of Pittsburgh
- University of Southern California
- University of Texas at Austin
- University of Texas at El Paso
- University of Utah
- University of Washington
- University of Wyoming

Q1 14



Charged Usage by Discipline




System Usage and Availability



Scheduled Availability



XDMOD Usage April 2013-March 2014



Blue Waters 16 FP "core equivalents" per node


Activity		Jobs	CPU Time (h)		Wait Time (h)	Wall Time (h)		Processors	
Users:	Pls:	Total:	Total:	Avg (Per Job):	Avg (Per Job):	Total:	Avg (Per Job):	Max:	Avg (Per Job):
541	113	364,183	4,954,274,681.9	13,603.81	9.28	1,199,070.5	3.29	859,648	5,993

XSEDE Total in "XSEDE SUs"

Activity				Jobs		Service (XD SU)		CPU Time (h)		Wait Time (h)	Wall Time (h)		Processors	
Users:	Pls:	Allocations:	Institutions:	Total:	Gateway:	Total:	Avg (Per Job):	Total:	Avg (Per Job):	Avg (Per Job):	Total:	Avg (Per Job):	Max:	Avg (Per Job):
5,193	1,756	2,004	567	5,652,588	212,035	5,423,072,922.0	959.40	1,764,664,787.6	312.19	3.26	15,623,218.6	2.76	98,304	112

Q1 14 Utilization

Rough Allocation Comparison in “core*hours”
Blue Waters - 7,466,160,704
INCITE 2013 - 4,681,410,000
INCITE 2014 - 5,780,000,000



Teams have Opportunities for more use

Auto-Refresh Include Backlog in Mouseover Tooltips

Start: 05/10/2014 12:11 pm

End: 05/11/2014 11:56 am

Queue Time Analysis is experimental at best, verify with command line

Overall Utilization Average: 70.41%


Node-Hours lost from Undersubmitted Workload (and iteration schedules): 64916.331 nodehours (12.07% avg)

Node-Hours lost by Down nodes: 351.195 nodehours (0.07% avg)

Node-Hours spent Draining: 93870.407 nodehours (17.45% avg)

Current Chart Resolution: 5 Minutes

2 people currently viewing this c



Lower priority jobs
tailored for backfill
will probably get
access often

Some time there is
not runnable work

Generalizations from PI interviews

- All PIs express a high satisfaction with Blue Waters and its services.
- About 1/5 of PIs expressed areas that they would like to see improved
- Many teams indicated they were getting close to publishing something significant based on Blue Waters use
- Some teams indicated are ready for visualization assistance
- Most of the teams were highly complementary of Blue Waters responsiveness and how quickly their jobs run
 - This is especially the case for those projects that are running very large scale jobs
- More than 1/2 the teams are using observed/experimental data as well as simulation
- A significant proportional of teams indicated that they were share the data of their results with others as well as papers.
- Several other teams indicated they would like to share their data and/or for the BW runs data might now have enough value to share with others. These observations led to our Data Sharing Services

SOME PLANNED ANNOUNCEMENTS

Pending concurrence from NSF

Blue Waters Data Sharing Service Prototype (BW-DSS)

- Provides means to share BW S&E Team data sets with the scientific community
- Data can be shared directly from Lustre or Nearline
 - Shared from ~/share
 - Read only sharing for greater security
- Globus Online used for sharing of large data sets (implemented and tested)
 - Share with colleagues through the GO Blue Waters portal interface
 - PIs decide who can see project data sets
 - Globus Plus user plans provided by the Blue Waters project (soon)

Blue Waters Data Sharing Service Prototype (BW-DSS) ... cont'd

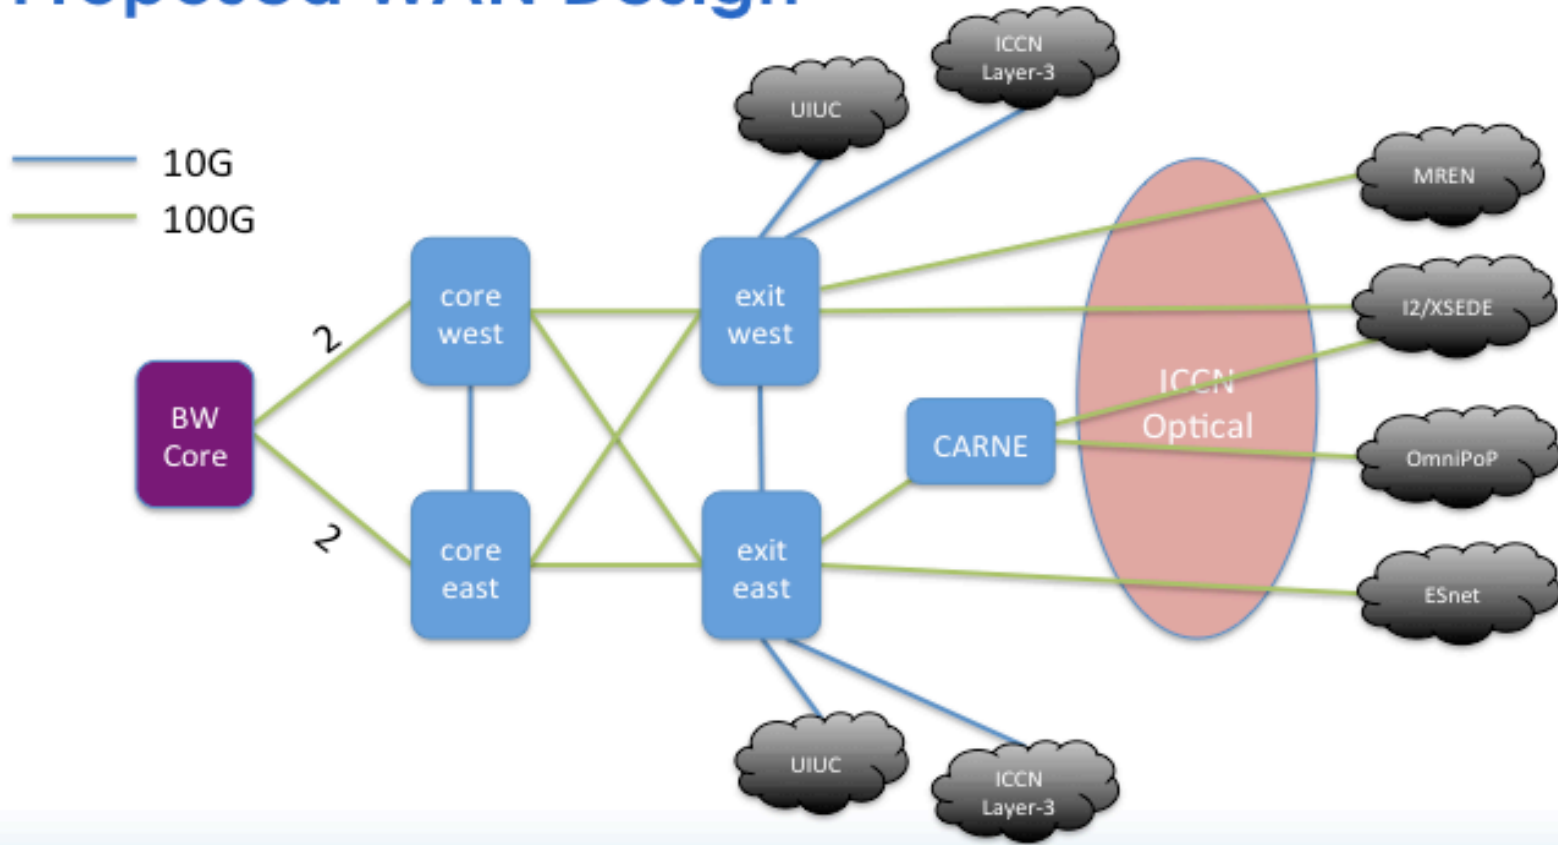
- Coming: Web portal interface for browsing and small data transfers
- Sharing enabled from `~/share/public_html/` on Lustre
 - Provides public read-only access
 - Can be used in conjunction with Globus Online sharing
- Preparation for tying into the National Data Service (NDS)
 - NFS proposal for scientific data retention and sharing on a national scale
 - NDS will make use of similar technologies like Globus Online
- Procedures will be defined – e.g. DOI and annotation requirements for sharing data
- Talk to Jason Alt, Galen Arnold or Mark Klein for more information

Innovative Data and Analysis Challenge

- NCSA and Cray have implemented a basic Mapreduce Framework running on Blue Waters
 - MR and a HDFS API on Lustre
 - Some limited tools
- Will provide up to 2 M node*hours total for teams wishing to try Blue Waters for innovative Data and Analysis projects
 - Light weight and fast proposal process outside NSF PRAC (similar to education allocations)
 - 6 month duration of projects
- Talk to Kalyana Chadalavada for more information

Possible Network Improvements

Proposed WAN Design



BLUE WATERS

SUSTAINED PETASCALE COMPUTING




GREAT LAKES CONSORTIUM
FOR PETASCALE COMPUTATION



Acknowledgements

This research is part of the Blue Waters sustained petascale computing project, which is supported by the National Science Foundation (award number OCI 07-25070) and the state of Illinois. Blue Waters is a joint effort of the University of Illinois at Urbana-Champaign, its National Center for Supercomputing Applications, and the Great Lakes Consortium for Petascale Computation.



Charles Seife

- Professor of journalism at NYU's Arthur L. Carter Journalism Institute
- Writes about physics and mathematics include 5 award winning books
 - Zero: The Biography of a Dangerous Idea (2000);
 - Alpha & Omega: The Search for the Beginning and End of the Universe (2003);
 - Decoding the Universe: How the New Science of Information is Explaining Everything in the Cosmos, From Our Brains to Black Holes (2005);
 - Sun in a Bottle: The Strange History of Fusion and the Science of Wishful Thinking (2008);
 - The Dark Arts of Mathematical Deception (2010).
- Writer for *Science*; *New Scientist*; *The Economist*; *Scientific American*; *The Philadelphia Inquirer*; *Discover*; *Slate*; *Smithsonian*; *The Washington Post*; *The New York Times*, and television documentaries about science and mathematics.
- Degrees in mathematics journalism from Princeton University, Yale University, Columbia University.

Getting Started with Data Sharing

- Fill out form from portal
 - <https://bluewaters.ncsa.illinois.edu/data-sharing>
- Get a Globus Online Plus Plan
 - <https://www.globus.org/account/GetPlus>
 - \$70/year
 - If you already have a Plus account from another institution, you are all set
 - Soon: Blue Waters project will sponsor your Plus plan

Globus Online (GO) sharing setup checklist

Paths in data space(s) shared are contained in :

__ ~/share/ (from GO endpoint: ncsa#BlueWaters , lustre filesystem)

__ ~/share/ (from GO endpoint: ncsa#Nearline , archive filesystem)

For each directory and dataset shared, metadata are provided as :

__README.txt __HDF_data __NetCDF_data

The person responsible for your team's shared data may be contacted via:

<email_address@your_institution.edu>

- 1) Add a [Globus Plus](#) subscription to your GO account.
- 2) See the [GO sharing setup guide](#) for assistance with setting up your endpoint

public_html sharing setup checklist

Public data may be shared from ~/share/public_html/ via http access. You may publicly share a URL of the form

[http://bluwaters.ncsa.illinois.edu/~\[user\]/share/public_html](http://bluwaters.ncsa.illinois.edu/~[user]/share/public_html). The portion of the path in **[user]** should correspond to your \$HOME on the system.

Your ~/share/public_html contains one of the files shown describing your data:

__index.html __README.txt __other (describe)



Manage Data

Groups

Support

jasonalt

[Update Profile](#)

[Change Password](#)

[Account Privacy](#)

[Manage Identities](#)

[Subscriptions](#)

[Log Out](#)



Subscribe to a Globus Plus Plan

[Select Plan](#) > [Payment](#) > [Confirm Order](#)

You currently subscribe to Globus Plus, [click here to manage your subscription.](#)

Request a Plus subscription from:

- GGenomics Plus Sponsor
- Globus Plus Sponsor
- KBase Plus Sponsor
- Michigan State University - Institute for Cyber-Enabled Research Plus Sponsor
- NCAR Plus Sponsor
- NIH Helix Systems Plus
- Penn State University Plus Sponsor
- Purdue University, RCAC Plus Sponsor
- San Diego Supercomputer Center
- The University of Missouri Research Network Plus Sponsor
- UNL Holland Computing Center
- University of Chicago Research Computing Center
- University of Chicago Vault Research Storage
- University of Colorado Boulder Research Computing Plus Sponsor
- Yale Biomedical HPC Center Plus Sponsor
- Yale Faculty of Arts & Sciences HPC Core Plus Sponsor



Manage Data

Groups

Support

jasonalt

[Transfer Files](#) | [Activity](#) | [Manage Endpoints](#) | [Dashboard](#) | [Flight Control](#)

Transfer Files

Get Globus Connect Personal
Turn your computer into an endpoint.

The screenshot displays the Globus Transfer Files interface with two side-by-side endpoint views. The left view is for endpoint `ncsa#BlueWaters` with path `/`, showing a list of system folders such as `bin`, `boot`, `cm`, `dev`, `etc`, `home`, `lib`, `lib64`, `local`, `media`, `mnt`, `opt`, `proc`, `projects`, `root`, `sbin`, `scratch`, `selinux`, `srv`, and `sw`. The right view is for endpoint `ncsa#Nearline` with path `/-/`, showing a list of folders including `Archive`, `no_quota`, `project.fy`, and `share`. A context menu is open over the `share` folder, listing options: `new folder`, `show hidden files`, `delete selected files`, and `share`. At the bottom, there is a section for 'Label This Transfer' with a text input field and a note: 'This will be displayed in your transfer activity.'

Transfer Files

Get Globus Connect Personal

Turn your computer into an endpoint.

The screenshot displays the Globus Transfer Files interface. At the top, there are two endpoint selection boxes. The left box is for 'nca#BlueWaters' with a path of '/'. The right box is for 'nca#Nearline' with a path of '/~'. Below these are two file browser panes. The left pane shows a standard Linux directory structure (bin, boot, cm, dev, etc, home, lib, lib64, local, med, mnt, opt, proc, proje, root, sbin, scra, selir, srv, sw). The right pane shows a directory structure with folders: Archive, no_quota, project.fy, and share. A modal dialog box titled 'Manage Shared Endpoint' is open in the foreground. It contains the following fields:

- Create New Shared Endpoint**
- Source Endpoint: nca#Nearline
- Source Path: /~/share/
- New Endpoint Name: jasonalt#share
- Description: This is how I share data from nearline

At the bottom of the dialog are two buttons: 'Create and Manage Access' and 'Cancel'.

Transfer Files

Get Globus Connect Personal
Turn your computer into an endpoint.

Manage Shared Endpoint

« shared endpoints list

Manage Permissions For jasonalt#share

Host: ncsa#Nearline:/-/share/

name	read	write
Path:/	view link for sharing	
Jason Alt (jasonalt)	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Galen Arnold (arnoldg)	<input checked="" type="checkbox"/>	<input type="checkbox"/>

✓1 user added successfully

ID (User or Group) search »

Path

NOTE: All paths are assumed to be folders

Permissions read write