

May 18, 2012

XSEDE for iPlant Users

John Fonner

Texas Advanced Computing Center

XSEDE

Extreme Science and Engineering
Discovery Environment



Outline

- What is XSEDE?
- What resources does XSEDE have?
- How does the XSEDE/iPlant connection work?
- How do I get my own allocation?



What is XSEDE?

- The Extreme Science and Engineering Discovery Environment (XSEDE): The most powerful integrated advanced digital resources and services in the world. Funded by NSF.
- A single virtual system that scientists can use to interactively share computing resources, data, and expertise.
- 9 supercomputers, 3 visualization systems, and 9 storage systems provided by 16 partner institutions (Service Providers or SPs)



What is XSEDE?

- The successor to the TeraGrid, XSEDE is an NSF-funded, advanced, nationally distributed open cyberinfrastructure, consisting of:
 - Supercomputing (and other computing)
 - Storage
 - Visualization
 - Data Collections
 - Network
 - Science Gateways
 - Unified Policies and Programs



XSEDE Service Providers

- NCSA, Illinois
- PSC, Pitt/Carnegie Mellon
- NICS, Tennessee/ORNL
- TACC, Texas
- SDSC, UC San Diego
- OSC, Ohio State Cornell Virginia
- Indiana Purdue Rice
- Shodor Foundation Argonne UC-Berkeley
- U Chicago SURA Open Science Grid



Allocation of XSEDE Resources

- XSEDE resources are allocated through a peer-reviewed process.
- Open to any US open science researcher (or collaborators of US researchers) regardless of funding source.
- XSEDE resources are provided at NO COST to the end user through NSF funding (~\$100M/year).



XSEDE Resources

XSEDE USER PORTAL
 Extreme Science and Engineering Discovery Environment

Search XSEDE...

SIGN IN

MY XSEDE | **RESOURCES** | DOCUMENTATION | ALLOCATIONS | TRAINING | USER FORUMS | HELP | ABOUT

Resource Monitor | Scheduled Downtimes | Queue Prediction | Remote Visualization | Science Gateways

HPC SYSTEMS | ADVANCED VIS SYSTEMS | STORAGE SYSTEMS | SPECIAL PURPOSE SYSTEMS

NAME	INSTITUTION	SYSTEM	PEAK TFLOPS	MEMORY TBYTES	STATUS	LOAD	RUNNING JOBS	QUEUED JOBS	OTHER JOBS
Kraken	NICS	Cray XT5	1174.00	147.00	Up	<div style="width: 100%;"></div>	245	855	392
Ranger	TACC	Sun Constellation Cluster	579.40	123.00	Up	<div style="width: 100%;"></div>	396	67	105
Lonestar	TACC	Dell Linux Cluster	302.00	45.00	Up	<div style="width: 100%;"></div>	183	588	47
Forge	NCSA	PowerEdge C6145 with NVIDIA Fermi M2070	150.00	64	Up	<div style="width: 100%;"></div>	11	19	0
Trestles	SDSC	Appro AMD Magny-Cours Cluster	100.00	20.25	Up	<div style="width: 100%;"></div>	1025	830	81
Steele	Purdue	Dell Intel 64 Linux Cluster	60.00	12.40	Up*	<div style="width: 100%;"></div>	1192	1805	75
Queen Bee	LONI	Dell Intel 64 Linux Cluster	50.70	5.31	Up	<div style="width: 100%;"></div>	24	0	0
Blacklight	PSC	SGI UV	37.20	32.00	Up	<div style="width: 100%;"></div>	54	200	9
Total:			2453.3	448.96			3130	4364	709

*Indicates failure of one or more status test.
 Hover mouse pointer over Resource Name, Resource Status, and headings to see additional information.

HPC SYSTEMS | ADVANCED VIS SYSTEMS | STORAGE SYSTEMS | SPECIAL PURPOSE SYSTEMS

NAME	INSTITUTION	SYSTEM	CPU	PEAK TFLOPS	MEMORY TBYTES	GRAPHICS HARDWARE
Longhorn	TACC	DELL/NVIDIA Visualization & Data Analysis Cluster	2048	20.70	13.50	128 NVIDIA Quadro Plex S4s NVIDIA FX 5800s
Nautilus	NICS	SGI/NVIDIA, Visualization and Data Analysis System	1024	8.19	4.00	16 NVidia graphics processing
Spur	TACC	Sun Visualization Cluster	128	1.13	1.00	4 NVIDIA Quadro Plex model Quadro Plex S4. Total: 32 NV GPUs
Total:			3200	30.02	18.5	

HPC SYSTEMS | ADVANCED VIS SYSTEMS | STORAGE SYSTEMS | SPECIAL PURPOSE SYSTEMS

NAME	INSTITUTION	SYSTEM	ONLINE STORAGE TB	OFFLINE STORAGE TB
Albedo Lustre-WAN	PSC	Lustre	1024	N/A
Data Capacitor	IU	Lustre	535	N/A
Data Replication Service	TACC	iRODS	N/A	1024
HPSS	NICS	HPSS	12	6220
NCSA Tape Storage	NCSA	EMC DiskXtender	250	10000
Ranch	TACC	Sun StorageTek Mass Storage Facility	110	40000
Total:			1931	57244

I'll run through TACC resources as an example...



Ranger: World Class Supercomputing Capability!



Ranger System Summary

- **Peak Performance – 579.4 Teraflops**
 - 3,936 Sun four-socket blades
 - 15,744 AMD “Barcelona” processors
 - Quad-core, four flops/clock cycle
- **Total Memory - 123 Terabytes**
 - 2 GB/core, 32 GB/node
 - 123 TB/s aggregate bandwidth
- **Interconnect – 1 GB/s, 1.6-2.85 μ sec latency**
 - Sun Data Center Switches (2), InfiniBand, up to 3456 4x ports each
 - 7.8 TB/s backplane
- **What’s all this mean?**
 - **Ranger debuted as the fastest open science supercomputer in the world... now, at 4 years old, still one of the world’s top 25 systems.**

Ranger Usage

- Who uses Ranger?
 - A community of researchers from around the US (along with international collaborators)
 - More than 2200 allocated users
 - 450 individual research projects
- Usage to Date?
 - >2,000,000 jobs have been run through the queues
 - >1 billion CPU hours consumed

Lonestar 4



Dell Intel 64-bit Xeon Linux Cluster
22,656 CPU cores (302 TFlops)
44 TB memory, 2.4 PB disk

Lonestar: The Stats

(We'll talk about what all this means later)

- 302 TeraFlops (Trillion floating point operations per second).
- 22,656 Intel Westmere Processor cores @3.33 Ghz in 1,888 Dell Blades.
- 2GB RAM per core, 24GB per node, **44.3 TB** total.
- 40Gbps QDR Infiniband interconnect in a fully non-blocking fat tree (Mellanox).
- 1.2Petabytes of disk in parallel filesystem (DataDirect, Lustre)

Specialized Subsystems

Lonestar is a comprehensive system for science and engineering

- High Performance Computing
- High Throughput Computing
- Visualization subsystem
- Large memory subsystem
- GPU/Accelerated Computing Subsystem

Stampede - Headlines

- 10PF+ Peak performance in initial system (2013)
 - 2PF conventional cluster
 - 8PF accelerated co-processor system
- 15PF+ after upgrade (2015)
- 14PB+ disk
- 200TB+ RAM
- 56Gb FDR InfiniBand interconnect
- Nearly 200 racks of compute hardware
- Integrated shared memory and remote visualization

Stampede

- Base Cluster:
 - Intel Sandy Bridge processors
 - Dell dual-socket nodes w/32GB RAM
 - More than 6,000 nodes
 - More than 100,000 cores
- Co-Processors:
 - Intel “MIC” Many Integrated Core processors
 - Special release of “Knight’s Corner” (>50 cores)
 - “Knight’s Ferry” development platforms at TACC now
- Total Concurrency approaching 500,000 cores

Subsystems

- Stampede will include 16 1TB Sandy Bridge shared memory nodes with dual GPUs.
- 128 of the compute nodes will be equipped with NVIDIA Kepler-class GPUs
- Storage subsystem will be Dell storage nodes.
 - Aggregate Bandwidth greater than 150GB/s
 - More than 14PB of capacity
 - Similar partitioning of disk space into two filesystems as Lonestar.

Longhorn: Most Powerful Interactive Remote Vis System in the World

- NSF “XD Visualization” award to TACC for \$7M to
 - deploy world-class visualization services for US open science research
 - also offers tremendous GPU computing capability
 - provide advanced support for 3 years
- **Longhorn specs**
 - 256 Dell Quad-core Intel Nehalem Nodes
 - 8 cores/nodes, 2048 total cores
 - 129 Nvidia Quadplexes
 - 4GPUs/node, 512 total GPUs)
 - QDR InfiniBand Interconnect

Stallion - Highest Resolution Display Environment in the World




Massive Computing Requires Massive Data Storage: Meet Corral!

- ~6 PetaByte DataDirect Networks online disk storage
 - 1 Petabyte = ~1 quadrillion characters, or 1 million gigabytes.
- Multiple access mechanisms
 - MySQL & Postgres SQL databases
 - Lustre parallel filesystem
 - iRODS
 - Web-based access
- ***Through Corral, TACC makes large datasets available to the community.***



Ranch Archival System

- Sun StorageTek Silo
 - 20,000 tapes
 - 100 PB capacity
 - Used for long-term storage



Other XSEDE Compute Resources

- Blacklight – PSC – Large scale shared memory (16TB)
- Kraken – NICS – 1PF Cray (>100k processors)
- (New) Gordon - SDSC – 350TF, Flash disk for data intensive computing
- Trestles – SDSC – Linux Cluster tuned for short jobs
- (Coming soon) Keeneland – NICS – GPU
- Forge – NCSA – GPU
- Queen Bee, Steele - ~50TF Linux Clusters



Other XSEDE specialized resources

- Visualization and Data Analysis
 - Nautilus
 - Longhorn, Spur
- Storage
 - Ranch, Data Capacitor, Albedo, HPSS, NCSA Tape
- Special
 - Condor Pool
 - OSG



iPlant and XSEDE

- When you run certain larger tools, through the Discovery Environment (e.g. Abyss assembler), you are already using XSEDE resources.
- You can also request direct, command line access to the large resources through the iPlant allocation (currently, Ranger, Lonestar, Blacklight, Trestles, Longhorn).
- Programmers using the iPlant API can also make calls to run directly on XSEDE systems.



iPlant: Ultra-High Throughput Inquiry



High Speed Data Intake



Abyss

TACC Lonestar

TACC Ranger


The screenshot shows the iPlant Collaborative Discovery Environment interface. It features a navigation sidebar on the left with icons for Data, Analysis, and Apps. The main workspace is divided into several panes: 'Data' (listing various datasets like 'vaughn', 'alee-s', 'analysis', etc.), 'Apps' (listing various applications like 'Workflows (26)', 'Public Applications (4)', etc.), and 'Analysis' (showing a table of analysis jobs with columns for Name, App, Start Date, End Date, and Status).

Name	App	Start Date	End Date	Status
PeakR	Workflow_demo	2012 Jan 25 13:30...	2012 Jan 25 13:34...	Completed
Extract	Interproscan_Hybrid...	2012 Jan 25 13:27...	2012 Jan 25 13:27...	Running
SAMTo	Checktrans	2012 Jan 25 12:05...	2012 Jan 25 12:04...	Completed
PeakR	Translat_jim	2012 Jan 25 12:05...	2012 Jan 25 12:54...	Completed
SAMTo	Translat_fir_ji_lex	2012 Jan 25 12:45...	2012 Jan 25 12:48...	Completed
EMBOSS	Translat_3hd	2012 Jan 25 12:34...	2012 Jan 25 12:37...	Completed
EMBOSS	Hybrid_Assembly_3hd...	2012 Jan 25 12:10...	2012 Jan 25 12:14...	Completed
Minimus	Trinity_scep_Merge-2...	2012 Jan 25 11:51...	2012 Jan 25 12:07...	Completed
TrinityR	Trinity_scep_metrics	2012 Jan 25 11:47...	2012 Jan 25 11:52...	Completed
TrinityR	Trinity_abyss_metrics	2012 Jan 25 11:47...	2012 Jan 25 11:50...	Completed
Translat	Merge_mn_Trinity_ah...	2012 Jan 25 10:46...	2012 Jan 25 10:51...	Completed

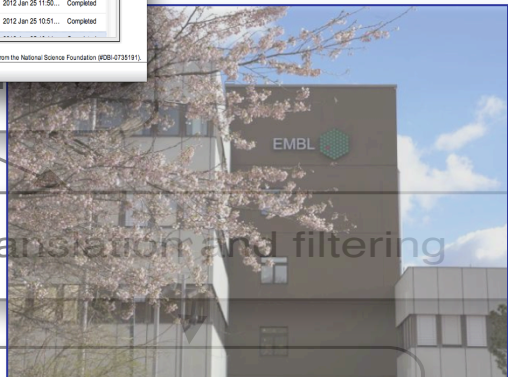


Contig Assessment

BLASTX Assessment



TACC Corral



Transition and filtering

InterproScan

PSC Blacklight

EBI Web Services



iPlant and XSEDE

- Through any of these methods, there are limits on how much of each resource iPlant can supply.
 - iPlant users used ~2 million CPU hours last year on XSEDE resources.
 - A similar amount will be available this year.
- When you reach a certain usage level, you will need to apply for your own XSEDE Allocation.
 - Peer reviewed, but open to any US researcher.



How do I get started using XSEDE

- To get started using XSEDE a researcher needs to:
 - Apply for an Allocation, or
 - Get added to an existing allocation
- To do either of these things, you should start with the **XSEDE User Portal**



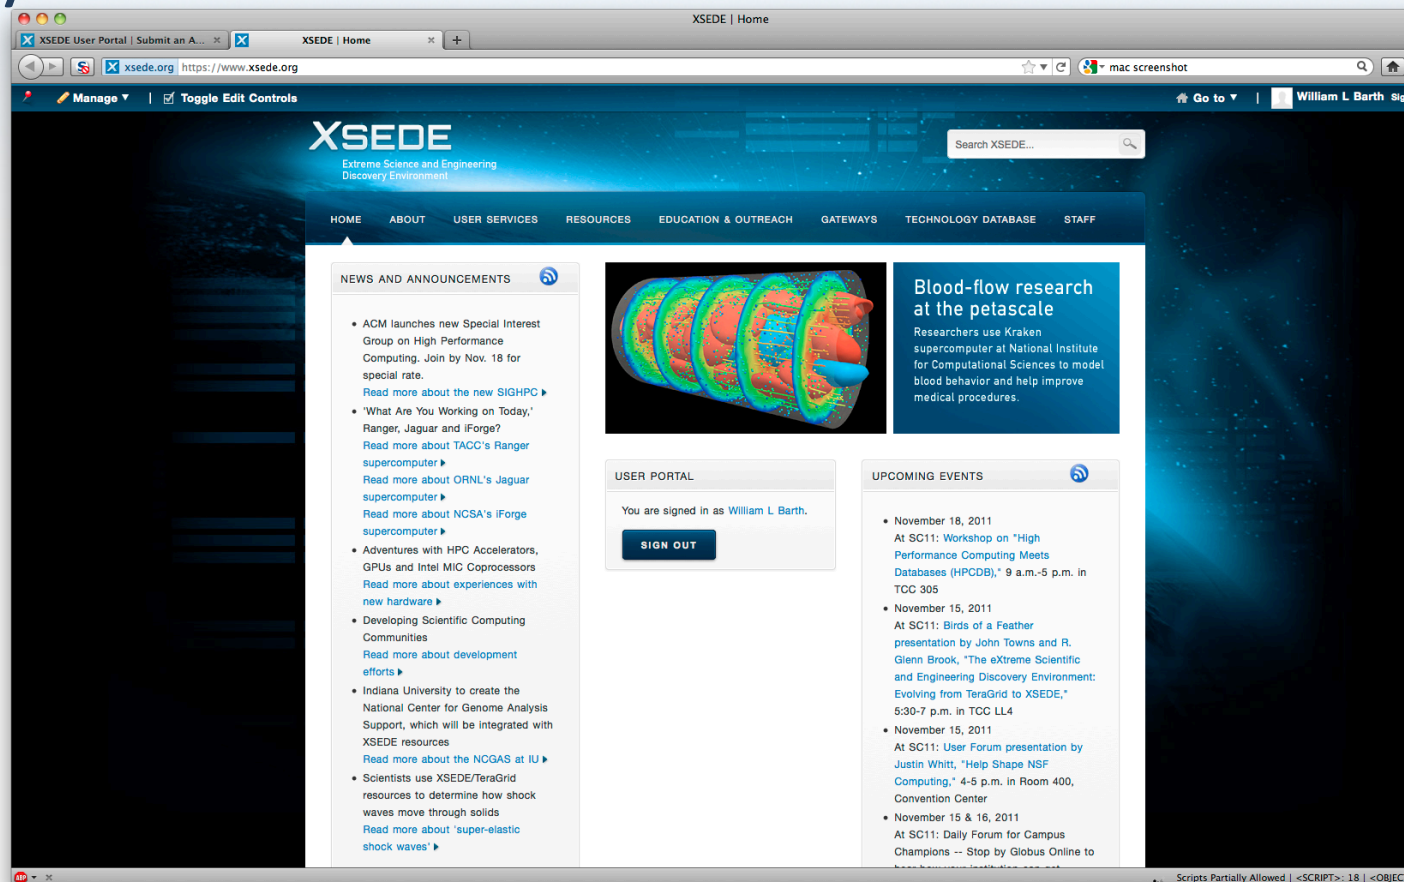
XSEDE User Portal (XUP)

- Web-based single point of contact that provides:
 - Continually updated information about your accounts.
 - Access to your XSEDE accounts and allocated resources:
 - A single location from which to access XSEDE resources. One can access all accounts on various machines from the Portal.
 - Interfaces for data management, data collections, and other user tasks and resources
 - Access to the Help Desk



The XSEDE.org Home Page

- From here, you can create a web account at any time!



User Responsibilities and Security

- The first time you login to the Portal, at the beginning of each allocation term, you will be asked to accept the User Responsibilities form:
 - Explains acceptable use to protect shared resources and intellectual property. Acknowledgment in publications, etc.
 - You are responsible for your account: Do not share accounts
 - User is responsible for protecting the passwords:
 - Includes not sharing passwords, not writing passwords down where they can be easily found, and not using tools which expose passwords on the network
 - This includes private keys: make sure they are password-protected.
- Appropriate Behavior
 - Protecting computing, closing SSH terminals when done, logging out of the User Portal when done, etc.
- Report Suspicious Activity.
- *If you have any suspicion that your account or personal computer has been compromised send email to help@xsede.org or call 24/7 1-866-907-2383 immediately.*



Getting an Allocation

- If you do not yet have an allocation, you can use the portal to acquire one.
- If you are a first time investigator, request a startup allocation.

The screenshot shows the XSEDE User Portal interface. The main content area is titled "Welcome to POPS: System for XSEDE Allocation Requests". Below this, there is a "Click to Enter or View a Request" button. The page provides information on how to submit an allocation, including links to the Request Steps Guide and XSEDE How to get an Allocation. It also mentions that users can refer to the XSEDE Resource Allocation Policies document for procedures or policies. A "Submission Schedule" section follows, with a table detailing the schedule for different request types.

REQUEST	SERVICE UNITS (SUS) RANGE (K=1000)	OPEN SUBMISSIONS	CLOSE SUBMISSIONS	ALLOCATIONS BEGIN	REVIEW CYCLE	
New Startups (Not normally renewable)	Grand total limit of 200K See Hardware Resource Catalog for specific startup limits.	Year round	n/a	Usually 2-3 weeks after submitted	Year round	
	Grand total limit of 200K See Hardware Resource Catalog for specific educational limits.	Year round	n/a	Usually 2-3 weeks after submitted	Year round	
Educational (Renewable)						
Research	No SU limit		Dec. 15	Jan. 15	Apr. 1	Quarterly
			Mar. 15	Apr. 15	Jul. 1	
			Jun. 15	Jul. 15	Oct. 1	
			Sept. 15	Oct. 15	Jan. 1	

Resource Changes

- SDSC Dash, PSC People, NCSA Ember, LONI/QueenBee Resources Removed
- SDSC, New Trestles System, New Gordon Compute Cluster, New Gordon ION
- NCSA, Forge (Farm) GPU System to replace Lincoln.

More resource information can be found on the [Resources Overview page](#)



Creating an Allocation

XSEDE POPS
Extreme Science and Engineering
Discovery Environment

POPS
Submission Home
Contact POPS

Submission Home
User: Barth, William

Research Requests:

- New**
CCR100034 Submission Status: **Approved** Active Status: **Active**
Title: **Developing Advanced MPI capabilities for Petascale Applications**
Submitted On: **Jul 16, 2010** Award Start Date: **Oct 01, 2010** Award End Date: **Mar 30, 2012**
Actions: [[View](#)] [[Transfer](#)] [[Supplemental](#)]

Staff Requests:

- New**
STA110019 Submission Status: **Approved** Active Status: **Active**
Title: **XSEDE SP TACC**
Submitted On: **Sep 19, 2011** Award Start Date: **Sep 20, 2011** & Award End Date: **Sep 20, 2012**
Actions: [[View](#)] [[Transfer](#)] [[Supplemental](#)]

New Request: Select An Appropriate Category For A New Request

- **Research:** Next submission period will be 12/15/2011 - 01/15/2012
These Research requests will be for allocation period 04/01/2012 - 03/31/2013
- **Startup:** For investigators new to XSEDE. Accepted anytime.
For details on the request limits for Startup allocations, [please see the allocations policies.](#)
- **Educational:** Specifically for classroom instruction and training courses. Accepted anytime.
- **Campus Champions**



Once your allocation is approved:

- The PI (principle investigator), Co-PI, or Allocation Manager can add users to an existing allocation through the portal.
 - XSEDE UserPortal:My XSEDE->Add/Remove User
- Takes the portal name of the user you want to add/remove.
- Accounts at certain Service Providers need to be activated before they can be accessed.



Accessing XSEDE Resources

- Several methods are possible:
 - Direct login access
 - Single Sign On (SSO) through portal
 - SSO between resources
 - Through Science Gateways
- Your choice of method may vary with:
 - How many resources you use
 - How much you want to automate file transfers, job submission, etc.



Accessing Resources(2)

- SSO is the default method; you'll need to file a ticket to request a direct access password to the machine.
- Direct access:
 - Use a secure shell (ssh) client.
 - From Linux or Mac terminal window:
 - `ssh -l <username> <machinename>`
 - E.g: `ssh -l jfonner ranger.tacc.utexas.edu`
 - From Windows:
 - Download one of many ssh clients
 - Free ones include “putty” and “SSH Client 3.2.9”
 - Most campuses have a site license for a fancier one.



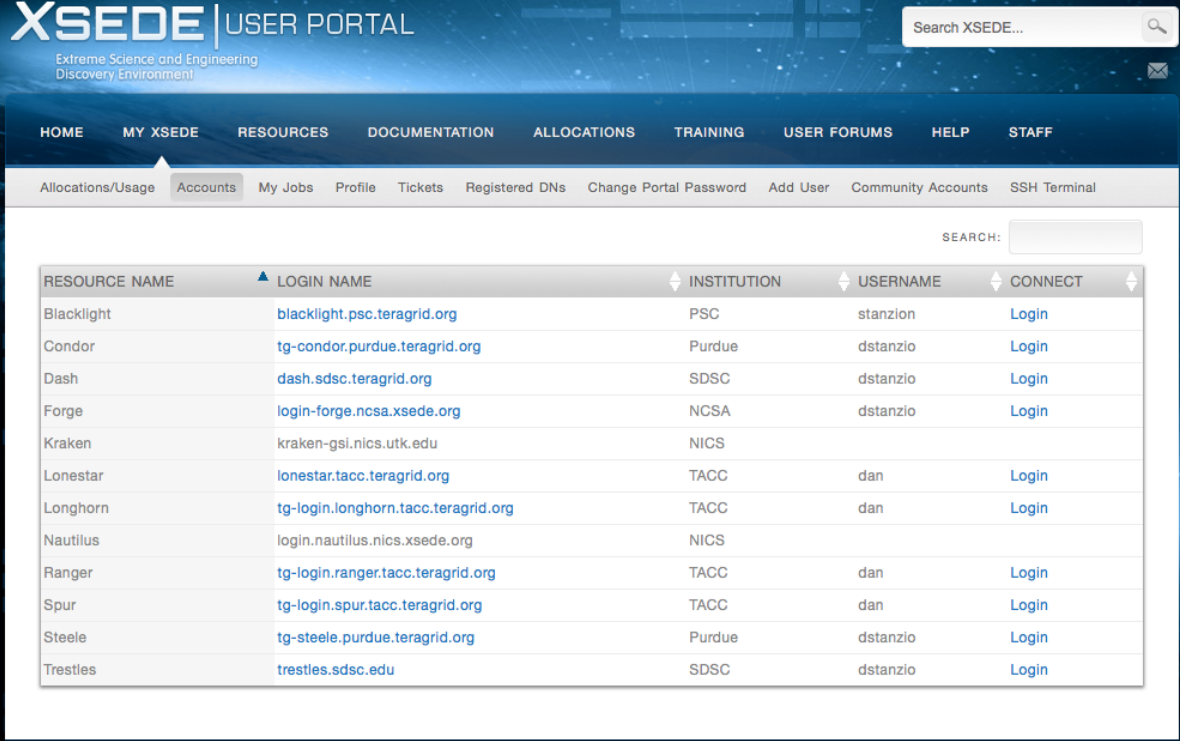
Single Sign On

- Single Sign-On (SSO) allows you to use just one username and password (your User Portal one) to log into every digital service on which you have an account.
- The easiest way to use SSO is via the XSEDE User Portal, but you can also use SSO via a desktop client or with an X.509 certificate.
- Web client:
<https://portal.xsede.org/group/xup/ssh-terminal>
- After you authenticate using SSO with your User Portal username and password, you will be recognized by all XSEDE services on which you have account, without having to enter your login information again for each resource.



SSO thru user portal

- Make sure you are logged into the XSEDE User Portal
- Go to 'My XSEDE' tab
- Go to the 'Accounts' link
- Resources you have access to will be indicated by a 'login' link
- Click on the 'login' link of the resource you would like to access.

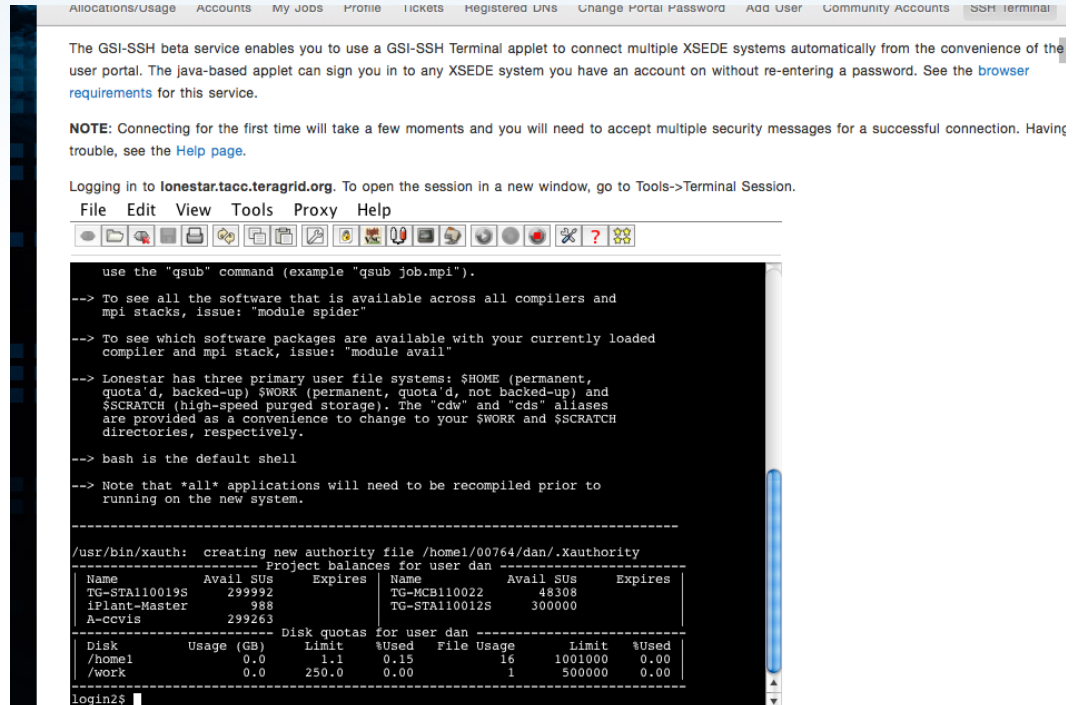


The screenshot shows the XSEDE User Portal interface. The header includes the XSEDE logo and the text 'Extreme Science and Engineering Discovery Environment'. A search bar is located in the top right corner. The main navigation menu includes 'HOME', 'MY XSEDE', 'RESOURCES', 'DOCUMENTATION', 'ALLOCATIONS', 'TRAINING', 'USER FORUMS', 'HELP', and 'STAFF'. Below this, a secondary menu shows 'Allocations/Usage', 'Accounts', 'My Jobs', 'Profile', 'Tickets', 'Registered DNs', 'Change Portal Password', 'Add User', 'Community Accounts', and 'SSH Terminal'. The 'Accounts' tab is active, displaying a table of resources with columns for Resource Name, Login Name, Institution, Username, and Connect. A search bar is also present above the table.

RESOURCE NAME	LOGIN NAME	INSTITUTION	USERNAME	CONNECT
Blacklight	blacklight.psc.teragrid.org	PSC	stanzion	Login
Condor	tg-condor.purdue.teragrid.org	Purdue	dstanzio	Login
Dash	dash.sdsc.teragrid.org	SDSC	dstanzio	Login
Forge	login-forge.ncsa.xsede.org	NCSA	dstanzio	Login
Kraken	kraken-gsi.nics.utk.edu	NICS		
Lonestar	lonestar.tacc.teragrid.org	TACC	dan	Login
Longhorn	tg-login.longhorn.tacc.teragrid.org	TACC	dan	Login
Nautilus	login.nautilus.nics.xsede.org	NICS		
Ranger	tg-login.ranger.tacc.teragrid.org	TACC	dan	Login
Spur	tg-login.spur.tacc.teragrid.org	TACC	dan	Login
Steele	tg-steele.purdue.teragrid.org	Purdue	dstanzio	Login
Trestles	trestles.sdsc.edu	SDSC	dstanzio	Login

SSO Thru User Portal

- A Java Applet will talk... you may be asked permission to allow it to run.
- After the applet starts, a blank terminal window will appear in your web browser.
- The window will fill with text indicating that you have been successfully logged into the resource of your choice.
- You can now work on this machine, and connect to other machines from this terminal, using the command *gssish machine-name*



The screenshot shows the XSEDE user portal interface. At the top, there are navigation links: Allocations/Usage, Accounts, My Jobs, Profile, Tickets, Registered LNs, Change Portal Password, Add User, Community Accounts, and SSH Terminal. Below the navigation is a paragraph explaining the GSI-SSH beta service and a link to browser requirements. A note indicates that first-time connections require security messages. The main content area shows the user logging in to `lonestar.tacc.teragrid.org`. Below the login message is a terminal window with a menu (File, Edit, View, Tools, Proxy, Help) and a toolbar. The terminal output includes instructions on using the `qsub` command, software availability information, file system details for `lonestar`, and a table of disk quotas for user `dan`.

```
----- Project Balances for user dan -----
Name      Avail SUS  Expires  Name      Avail SUS  Expires
TG-STAL10019S  299992                TG-MCB110022  48308
iPlant-Master  988                TG-STAL10012S  300000
A-cvvis      299263

----- Disk quotas for user dan -----
Disk  Usage (GB)  Limit  Used  File Usage  Limit  Used
/home1  0.0         1.1    0.15  16          1001000  0.00
/work  0.0         250.0  0.00  1           500000  0.00
```



Another Access Path: Science Gateways

- There are many sites that give you web-based, domain-specific access to applications running on XSEDE.
 - Collectively, we call them “Science Gateways”
 - View a list of them on the User Portal, in the “Resources” tab.
 - Access methods vary; click on the specific gateway to find out more (dozens available, across many fields!).
- iPlant’s DE is one of these gateways!



The Mobile User Portal



- <https://mobile.xsede.org>
- Allows browsing of all XSEDE systems, file downloading, and third-party transfers.
- It provides several features for mobile users such as one touch file publishing to the user's public folder, simple creation of shared groups for any file/folder, and one click permission management of all XSEDE systems, file downloading, and third-party transfers.

XUP Resource Monitor

- View system information: TFLOPS, memory, today's load, jobs running in queue. **Status: up or down:** takes you to the news announcements that tells you when the machine is expected to come back up.

NAME	INSTITUTION	SYSTEM	PEAK TFLOPS	MEMORY TBYTES	STATUS	LOAD	RUNNING JOBS	QUEUED JOBS	OTHER JOBS
Kraken	NICS	Cray XT5	1174.00	147.00	Up	<div style="width: 100%;"></div>	186	648	180
Ranger	TACC	Sun Constellation Cluster	579.40	123.00	Up	<div style="width: 100%;"></div>	369	75	112
Lonestar	TACC	Dell Linux Cluster	302.00	45.00	Up	<div style="width: 100%;"></div>	313	467	86
Forge	NCSA	PowerEdge C6145 with NVIDIA Fermi M2070	150.00	64	Up	<div style="width: 0%;"></div>	0	6	0
Trestles	SDSC	Appro AMD Magny-Cours Cluster	100.00	20.25	Up	<div style="width: 100%;"></div>	181	0	49
Steele	Purdue	Dell Intel 64 Linux Cluster	60.00	12.40	Up*	<div style="width: 100%;"></div>	2414	822	66
Queen Bee	LONI	Dell Intel 64 Linux Cluster	50.70	5.31	Up*	<div style="width: 100%;"></div>	40	25	0
Blacklight	PSC	SGI UV	37.20	32.00	Up	<div style="width: 100%;"></div>	71	11	15
Dash	SDSC	Appro Intel Nehalem Cluster	4.90	3.00	Up	<div style="width: 0%;"></div>	1	0	0
Total:			2458.2	451.96			3575	2054	508

**Indicates failure of one or more status test.
Hover mouse pointer over Resource Name, Resource Status, and headings to see additional information.*



User Portal: User Forums

- The User Forums are a great place to ask question, get help, or discuss ideas about XSEDE.

The screenshot shows the XSEDE User Portal interface. At the top, there is a navigation bar with links for HOME, MY XSEDE, RESOURCES, DOCUMENTATION, ALLOCATIONS, TRAINING, USER FORUMS, HELP, and STAFF. A search bar is located in the top right corner. Below the navigation bar, the page title is "Forums". A paragraph explains that the forums are for users to share experiences, questions, and comments, but visitors must be logged in to contribute. Below this, there are links for "Message Boards Home", "Recent Posts", "My Posts", "My Subscriptions", and "Statistics", along with a "SEARCH" button. There are also links for "RSS (Opens New Window)" and "Subscribe". A "POST NEW THREAD" button is prominently displayed. The "Categories" section is expanded, showing a table with columns for CATEGORY, CATEGORIES, THREADS, and POSTS. The categories listed are Allocations, Campus Champions, Everything Fortran, Everything OpenMP, and General Discussion. Below the table, it says "Showing 5 results." The "Threads" section is expanded, showing a message: "There are no threads in this category."

CATEGORY	CATEGORIES	THREADS	POSTS	
Allocations Discuss allocations questions regarding applying for an allocation on XSEDE	0	1	1	Actions
Campus Champions A forum for campus champions to share ideas, discuss issues, and collaborate.	0	2	6	Actions
Everything Fortran	0	0	0	Actions
Everything OpenMP	0	0	0	Actions
General Discussion Please post threads under this general XSEDE User Forum category, we will create sub-categories as relevant threads emerge.	0	15	47	Actions



Running Jobs

- Each system in XSEDE has some local options you will need to know about to run jobs.
- To learn about the specifics of each system, check out the user guides:
 - In the portal, under “Documentation” select “User Guides”
- Pay particular attention to:
 - File Systems
 - Batch job submission



File Systems on XSEDE Resources

- Where your data resides on XSEDE and the appropriate storage is your responsibility. In general, all resources provide:
 - HOME: Permanent space, but small. A good choice for building software and working file collections of small to medium sized files, where a medium sized file is less than 50 MB.
 - SCRATCH: More space, but temporary; use while you are running your jobs. *Scratch space is temporary; it is not backed up and has limited redundancy, and is periodically purged of old files!*
 - Archival Storage: Long term storage of large amounts of data (often tape); slower access, accessible from all sites.



Batch Jobs

- All XSEDE compute resources use some form of batch scheduler.
- Compute jobs **can not** be run on the login nodes (no faster than a normal workstation!)
- There are several batch systems in use, but all work basically the same way.
 - Request number/type of nodes you need.
 - Specify how long you need to run.
 - Specify where your output files go.
 - Jobs typically described with a job script:



XSEDE

Sample Job Script for Grid Engine on TACC Lonestar:

```
#!/bin/bash
```

```
#$ -N myMPI # Job Name
```

```
#$ -j y # Combine stderr and stdout
```

```
#$ -o $JOB_NAME.o$JOB_ID # Name of the output file
```

```
#$ -pe 12way 12 # Requests 12 tasks/node, 12 cores total
```

```
#$ -q normal # Queue name "normal"
```

```
#$ -l h_rt=01:30:00 # Run time (hh:mm:ss) - 1.5  
hours
```

```
./run -param1 1 -param2 "a" -flag -output foobaz.txt
```



XSEDE

Submitting/manipulating batch jobs

- Submit the script that you have created:
 - Actual commands are machine specific, but they follow general principles

```
qsub jobname
```

```
qstat -a
```

```
qstat -u username
```

```
qdel jobid
```

```
man qsub
```



Managing Your Environment: Modules

- Allows you to manipulate your environment.
- 'module list' shows currently loaded modules.
- 'module avail' shows available modules.
- 'module show' <name> describes module.

<http://modules.sourceforge.net/>

```
% module load gcc/3.1.1
% which gcc
/usr/local/gcc/3.1.1/linux/bin/gcc
% module switch gcc/3.1.1 gcc/3.2.0
% which gcc
/usr/local/gcc/3.2.0/linux/bin/gcc
% module unload gcc
% which gcc
gcc not found
```



ECSS – Extended Collaborative Support Services

- Expertise available in a wide range of areas
 - Performance analysis
 - Petascale optimization
 - Gateways and web portals
 - Specialized scientific software
- Can solicit ECSS support at any time through the Allocations tab in the XSEDE User Portal
- Requires written justification and a project plan
- Inquire at help@xsede.org



ECSS can include

- Porting applications to new resources
- Providing help for portal and gateway development
- Implementing algorithmic enhancements
- Implementing parallel math libraries
- Improving scalability of codes to higher processor counts
- Optimizing codes to efficiently utilize specific resources
- Assisting with visualization, workflow, data analysis, and data transfer



Reporting and Tracking Issues

- portal.xsede.org -> Help
 - Help Desk: Submit ticket
 - Security Incidents
 - E.g. your account has been compromised
- portal.xsede.org -> My XSEDE -> Tickets
 - Submit ticket
 - View past tickets (both open and closed)
- Email: help@xsede.org
- Phone: 1-866-907-2383 (24x7x365)



Questions? Need Help?

- First, try searching the knowledge base or other documentation
- Next, submit a ticket
 - portal.xsede.org -> My XSEDE -> Tickets
- Send email
 - help@xsede.org
- Or call the Help Desk
 - 1-866-907-2383



Need more training?

- portal.xsede.org -> Training
 - Course Calendar
 - On-line training



**Thanks for listening, and welcome to iPlant
and XSEDE!**

jfonner@tacc.utexas.edu




File Transfers: Small Files (<2GB)

- To transfer *small files* between XSEDE Resources and/or your own workstation you can use **scp** or **sftp**.
- From Linux or Mac, you can run these commands directly from the terminal
- From Windows, use your ssh client to do this (putty has free downloads for these tools, too! – just Google “putty sftp”).
- These are easy to use and secure, but provide poor performance for large files.



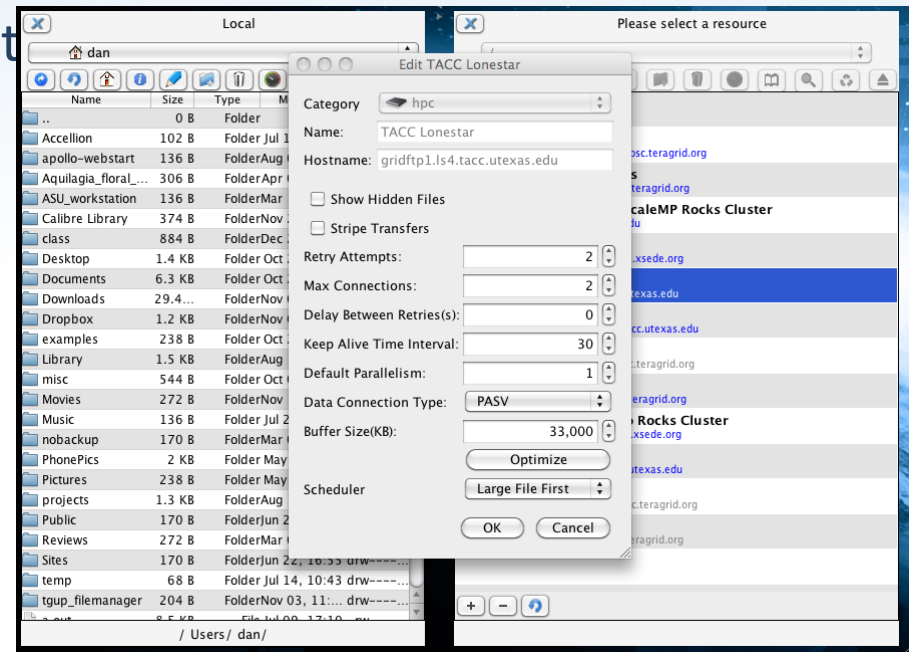
File Transfer: User Portal

- Log into the XSEDE User Portal
- Select “Resources” tab
- Select “File Manager” tab
 - (now wait for Java Applet to load)
 - May need to allow access for applet to run by clicking OK.
- You will see a list of all machines. *This includes:*
 - your local machine.
 - XSEDE\$Share: 2GB of space to collaborate. Allows you to share files with your collaborators.



Transferring Large Files with User Portal

- For large file transfers, we need to set a few parameters.
- Before clicking on the resource, Right click on the resource you're going to transfer data from and select Edit. This will bring up the file transfer parameters:
- Click the checkbox next to **“Stripe Transfers”** - Click OK
- Repeat for the other panel using the destination resource
- Repeat this every time you change Resources
- Drag and drop the file from source to destination to transfer.



What is Globus Online?

- Initial implementation of XSEDE User Access Services (XUAS)
- Reliable data movement service
 - High performance: Move terabytes of data in thousands of files
 - Automatic fault recovery
 - Across multiple security domains
- Designed for researchers
 - Easy “fire and forget” file transfers
 - No client software installation
 - New features automatically available
 - Consolidated support and troubleshooting
 - Works with existing GridFTP servers
 - Ability to move files to any machine (even your laptop) with ease




"We have been using Globus Online to move files to a TeraGrid cluster where we analyze and store tens of terabytes of data... I plan to continue using GO to access these resources within XSEDE to easily get my files where they need to go."
-- University of Washington user



"The service is reliable and easy to use, and I look forward to continuing to use it with XSEDE. I've also used the Globus Connect feature to move files from TeraGrid sites to other machines -- this is a very useful feature which I'm sure XSEDE users will want to take advantage of."
-- NCSA user



Case Study: Indiana University



- **Fast:** Reduced transfer times
- **Easy:** Fire-and-forget transfers
 - Automated retry
 - No file pre-staging
 - No complex infrastructure
 - Convenient CLI or GUI interfaces

Indiana University researcher moved ~6 TB from Oak Ridge to TACC in 2 days

TACC



“Globus Online frees up my time to do more creative work than typing scp commands or devising scripts to initiate and monitor progress to move many files.”

“I moved 100 7.3 GB files tonight in about 1.5 hours. I am very impressed. I also like the new commands and help system.”



Getting Started (2 easy steps)

1. **Sign up:** Visit www.globusonline.org to create an account

The image shows a screenshot of the Globus Online 'Create an account' page. The page has a blue header with the 'globus online' logo and a link for 'Already a member? Sign-in.'. The main form contains the following fields: 'Full Name' (with a placeholder 'First Last'), 'Email' (with 'float@institute.org'), 'Username' (with 'float'), 'Password' (with a strength indicator 'Strong'), and 'Confirm Password'. Below the form is a 'Terms of Service' section with a checkbox for 'I agree to the Terms of use.' and 'Cancel' and 'Register' buttons. A red arrow points to the 'Sign Up' button in a callout box. Another callout box shows a message: 'Please standby while your file transfer account is created.' with a loading spinner and the text 'Provisioning in Progress (This may take a few minutes.)'.

Getting Started (2 easy steps)

2. Start moving files: Pick your data and where you want to move it, then click to transfer

The screenshot shows the Globus Online interface. The top left panel, titled 'File Transfer', contains the text 'Use your browser to move data securely and reliably.' and 'Start Transfer' with a red arrow pointing to it. The main interface is split into two panes. The left pane, 'Transfer Files', shows a file browser for the source endpoint 'xsede#bigred'. The right pane shows a file browser for the destination endpoint 'smartin#ranger'. A blue callout box on the left points to the 'Endpoint' field in the left pane with the text 'File location: xsede endpoint'. A blue callout box on the right points to the 'Endpoint' field in the right pane with the text 'Destination server'. The top right panel, 'Transfer Activity', shows a table of transfer tasks with columns for Status, ID, Task Progress, Username, and Completion Time. A red arrow points to the 'View Transfers' dropdown menu in the top right corner.

Status	ID	Task Progress	Username	Completion Time
✓	5fd30...	3 / 3	vas	11/18/2010 07:31 PM
✓	f091a...	1 / 1	vas	11/18/2010 07:14 PM
✓	0793e...	1 / 1	vas	11/17/2010 08:55 PM
✓	049a3...	1 / 1	vas	11/17/2010 08:55 PM
✓	00d9d...	1 / 1	vas	11/17/2010 08:55 PM
✓	fd64...	1 / 1	vas	11/17/2010 08:55 PM

