

# BLUE WATERS

SUSTAINED PETASCALE COMPUTING

## Lustre Environment

BlueWaters




GREAT LAKES CONSORTIUM  
FOR PETASCALE COMPUTATION




UIUC/NCSA AND CRAY  
CONFIDENTIAL

Do not copy or distribute without expressed permission  
from the NCSA Blue Waters Project Office



Blue Waters 11Petaflop System

FDR IB



36 x Sonexion 6000

Lustre 2.1: > 25PB @ > 1TB/s

FDR IB



28 x Dell R720 IE nodes  
2 x 2.1GHz w/ 8 cores  
1 x 40GbE  
GridFTP access only

100 x 40GbE **HPSS**  
High Performance Storage System


440Gb/s



Internet

# Overall Environment

- Lustre file system > 1TB/sec aggregately
  - Lustre 2.1+ on servers; clients 1.8.6
  - 480 XIO nodes as Lnet routers to OSSs
    - 4:3 Fine grain routing
    - Lnet routers connect directly to top of rack Sonexion IB switches
  - Sonexion 1600s
    - All data currently striped RAID 6 8+2
    - 2 OSS per SSU



Sonexion  
1600 rack

# File Systems

- /scratch
  - 20 PiB usable space
  - 1440 OSTs / 180 SSU
  - 960 GiB/sec
- /home and /projects (1/10<sup>th</sup> /scratch)
  - 2 PiB usable space
  - 144 OSTs / 18 SSU
  - 96 GiB/sec

# File System Management

- Quotas (Lustre native)
  - 7 day grace period
  - /home 1TiB/1.5TiB Quota - enforced
  - /projects 5TiB/5.5TiB Quota - enforced
  - /scratch 500TiB/550TiB – not enforced
- Purge (/scratch)
  - Implemented in Robinhood by CEA
  - 30 day retention
  - Change log driven

# BLUE WATERS

SUSTAINED PETASCALE COMPUTING

## HPSS Nearline

BlueWaters

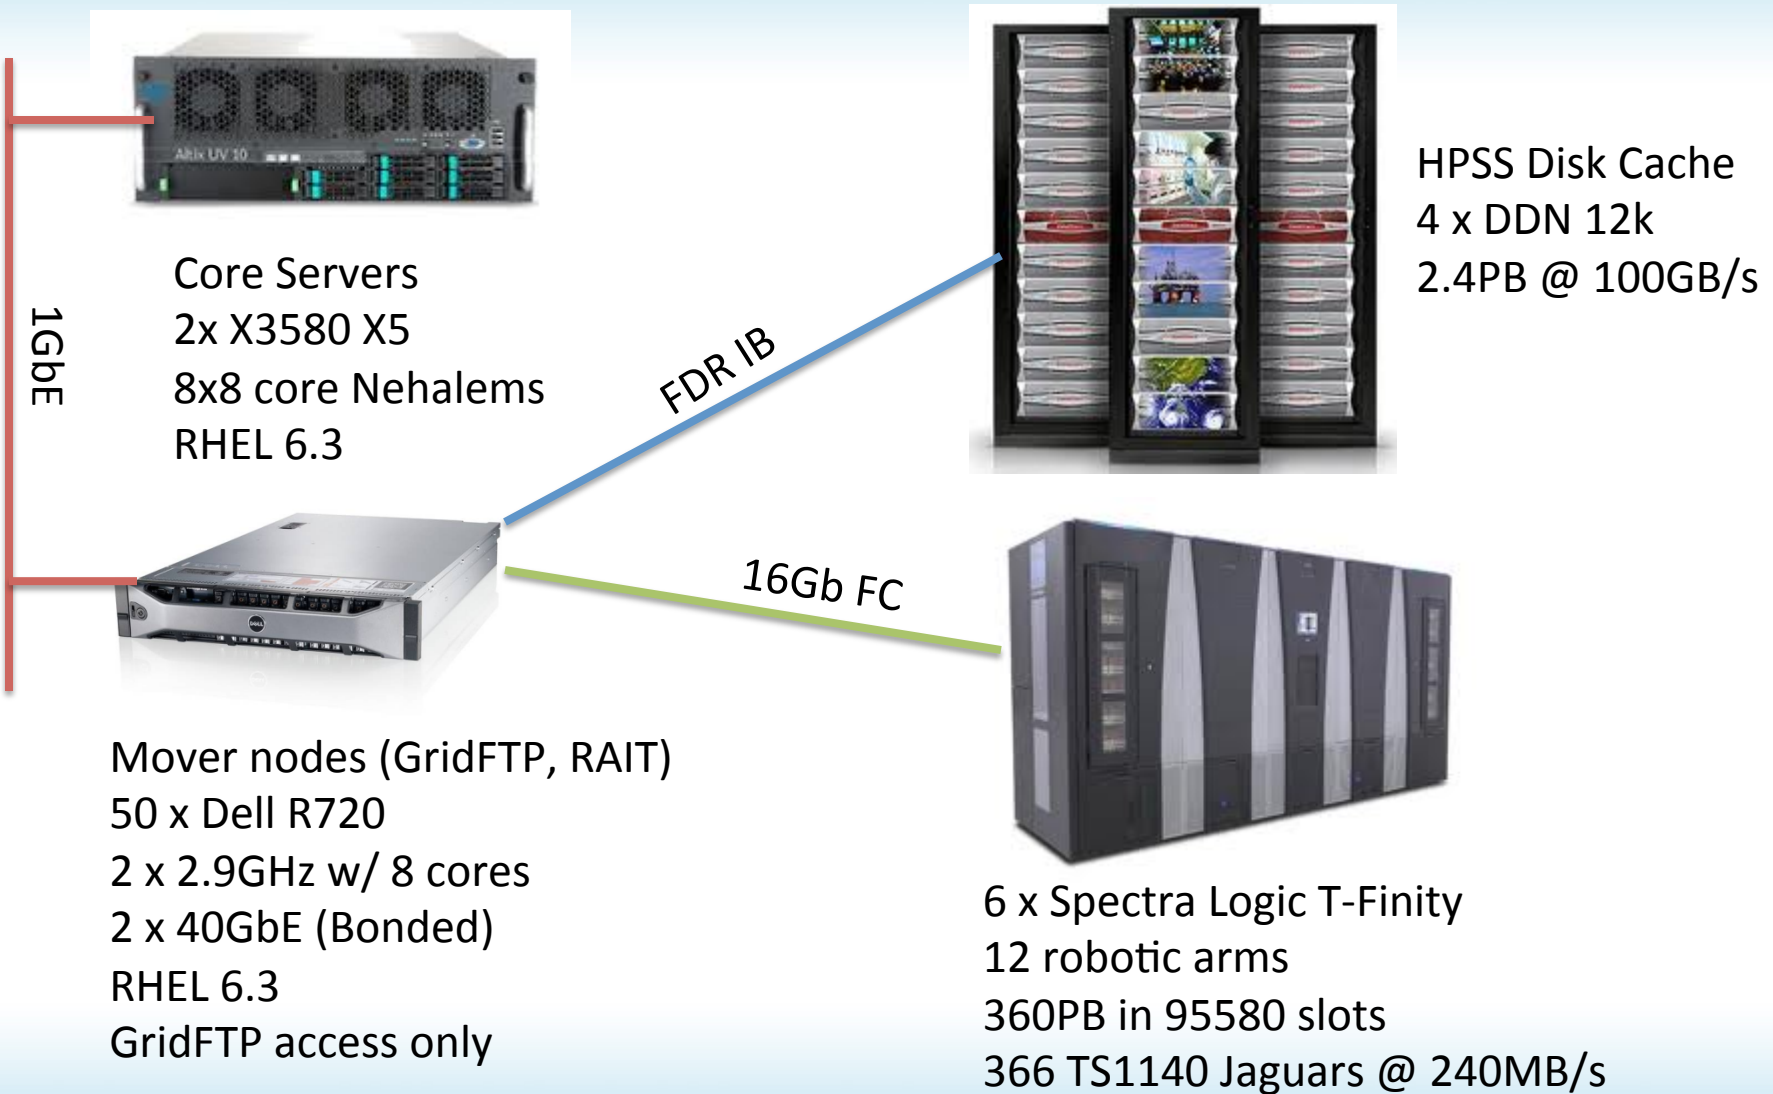


GREAT LAKES CONSORTIUM  
FOR PETASCALE COMPUTATION



UIUC/NCSA AND CRAY  
CONFIDENTIAL

Do not copy or distribute without expressed permission  
from the NCSA Blue Waters Project Office



## HPSS Directories you should care about

- Your home directory
  - /u/sciteam/bob (same as Lustre)
  - Default quota of 5TB. Can not be increased
- Your project directories
  - /projects/sciteam/jn0 (same as Lustre)
  - Default quota of 50TB. Can be increased with a request into the Blue Waters ticketing system
- No purge policy! Data stays for the life of your project



Globus Online | Reliable File Transfer

https://www.globusonline.org

g globus online

Manage Data | News & Events | About | Support | jasonalt

Reliable, high-performance, secure file transfer.  
**Move files fast. No IT required.**

**+ WATCH A VIDEO**  
Globus Online in a nutshell

**> GET STARTED**  
Sign up and get moving

**7,401,013,876 MB** TRANSFERRED

## This is how Blue Waters users 'see' HPSS

The screenshot shows the Start Transfer web interface. At the top, there are navigation links: "start transfer", "view transfer activity", "manage endpoints", and "dashboard". Below this, there are two transfer endpoints. The left endpoint is "jasonal#ncsa-mss" with a path of "/~/" and a file list. The right endpoint is "ncsa#forge" with a path of "/" and a file list. Both file lists include folders and files with their sizes. At the bottom, there is a "Label This Transfer" field with a placeholder text: "This will be displayed in your transfer activity."

Endpoint	Path	File Name	Size	Type
jasonal#ncsa-mss	/~	Archive		Folder
		DAR		Folder
		symlink		Folder
		test		Folder
		uberftp-client-2.6		Folder
		JUNK	10 GB	File
		LARGEFILE	4.29 GB	File
		PARTIAL	584 b	File
		decrypt	59 b	File
		gsi-cert-monitor	4.24 kB	File
		hpss-gridftp.tar.gz	5.18 kB	File
		hpss_krb5.keytab.tar.gz.enc	2.39 kB	File
		ie_krb5.keytab.tar.gz.enc	2.77 kB	File
		script	1.98 kB	File
		script.output	152.49 kB	File
		script.output2	82.08 kB	File
task-1.9.4.tar.gz	467.85 kB	File		
txt	55 b	File		
uberftp	567.24 kB	File		
uberftp.get	567.24 kB	File		
ncsa#forge	/	bin		Folder
		boot		Folder
		dev		Folder
		etc		Folder
		gpfs		Folder
		home		Folder
		ipaths		Folder
		lib		Folder
		lib64		Folder
		media		Folder
		mnt		Folder
		opt		Folder
		proc		Folder
		projects		Folder
		root		Folder
		sbin		Folder
scratch		Folder		
selinux		Folder		
srv		Folder		
sys		Folder		

View Transfers | Transfer | (x)

https://www.globusonline.org/xfer/ViewTransfers

globus online Manage Data | Support | jasonalt

start transfer | view transfer activity | manage endpoints | dashboard

## Transfer Activity

Cancel <<< 1 of 1 >>> View 25 Records

	Status	Label	Task Progress	Completion Time	Request Time
<input type="checkbox"/>	✓	Task Id:890b8f10-08ad-11e2-bf56-1231380b8963	1 / 1	09/27/2012 09:13 AM	09/27/2012 09:13 AM
<input type="checkbox"/>	✓	MSS to H2O	1 / 1	05/29/2012 10:13 AM	05/29/2012 10:12 AM
<input type="checkbox"/>	✓	H2O to MSS	1 / 1	05/29/2012 10:11 AM	05/29/2012 10:11 AM
<input type="checkbox"/>	✓	Task Id:5526357e-a99e-11e1-81e6-1231381bd061	1 / 1	05/29/2012 09:57 AM	05/29/2012 09:55 AM
<input type="checkbox"/>	✓	Task Id:40ef06cc-a99d-11e1-81e6-1231381bd061	1 / 1	05/29/2012 09:48 AM	05/29/2012 09:48 AM
<input type="checkbox"/>	✓	Task Id:3a564970-6c57-11e1-81e6-1231381bd061	1 / 1	03/13/2012 07:33 AM	03/12/2012 10:23 AM
<input type="checkbox"/>	✓	Task Id:b8a24d3c-c208-11e0-bc85-1231381a212f	1 / 1	08/08/2011 05:18 PM	08/08/2011 04:54 PM
<input type="checkbox"/>	✓	Task Id:bb4a295c-9778-11e0-b542-123139152864	1 / 1	06/15/2011 12:59 PM	06/15/2011 12:56 PM
<input type="checkbox"/>	✓	Task Id:a1e88bbc-9777-11e0-95fc-123139152864	1 / 1	06/15/2011 12:56 PM	06/15/2011 12:56 PM
<input type="checkbox"/>	✓	Task Id:3cb845e0-9775-11e0-bf0f-1231380e306c	1 / 1	06/15/2011 12:33 PM	06/15/2011 12:31 PM

Cancel <<< 1 of 1 >>> View 25 Records

# 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.

