



The Center for Computational Research & Grid Computing

Russ Miller

**Center for Computational Research
Computer Science & Engineering
SUNY-Buffalo**

Hauptman-Woodward Medical Inst

NSF, NIH, DOE
NIMA, NYS, HP



University at Buffalo
The State University of New York

Center for Computational Research 1999-2004 Snapshot

■ High-Performance Computing and High-End Visualization

- ❑ 110 Research Groups in 27 Depts
- ❑ 13 Local Companies
- ❑ 10 Local Institutions

■ External Funding

- ❑ \$116M External Funding
 - \$16M as lead
 - \$100M in support
- ❑ \$43M Vendor Donations
- ❑ Total Leveraged: \$0.5B

■ Deliverables

- ❑ 400+ Publications
- ❑ Software, Media, Algorithms, Consulting, Training, CPU Cycles...



Major Compute Resources

- **Dell Linux Cluster: #22® #25® #38® #95**
 - ❑ 600 P4 Processors (2.4 GHz)
 - ❑ 600 GB RAM; 40 TB Disk; Myrinet
- **SGI Origin3700 (Altix)**
 - ❑ 64 Processors (1.3GHz ITF2)
 - ❑ 256 GB RAM
 - ❑ 2.5 TB Disk
- **SGI Origin3800**
 - ❑ 64 Processors (400 MHz)
 - ❑ 32 GB RAM; 400 GB Disk
- **Apex Bioinformatics System**
 - ❑ Sun V880 (3), Sun 6800
 - ❑ Sun 280R (2)
 - ❑ Intel PIIIs
 - ❑ Sun 3960: 7 TB Disk Storage
- **HP/Compaq SAN**
 - ❑ 75 TB Disk
 - ❑ 190 TB Tape
 - ❑ 64 Alpha Processors (400 MHz)
 - ❑ 32 GB RAM; 400 GB Disk
- **Dell Linux Cluster: #187® #368® off**
 - ❑ 4036 Processors (PIII 1.2 GHz)
 - ❑ 2TB RAM; 160TB Disk; 16TB SAN
- **IBM BladeCenter Cluster: #106**
 - ❑ 532 P4 Processors (2.8 GHz)
 - ❑ 5TB SAN
- **IBM RS/6000 SP: 78 Processors**
- **Sun Cluster: 80 Processors**
- **SGI Intel Linux Cluster**
 - ❑ 150 PIII Processors (1 GHz)
 - ❑ Myrinet

CCR Visualization Resources

■ Fakespace ImmersaDesk R2

- ❑ Portable 3D Device

■ Tiled-Display Wall

- ❑ 20 NEC projectors: 15.7M pixels
- ❑ Screen is 11' ´ 7'
- ❑ Dell PCs with Myrinet2000

■ Access Grid Nodes (2)

- ❑ Group-to-Group Communication
- ❑ Commodity components

■ SGI Reality Center 3300W

- ❑ Dual Barco's on 8' ´ 4' screen



TOP500 List for June 2004

R_{max} and R_{peak} values are in GFlops. For more details about other fields, please click on the button "Explanation of the Fields"



EXPLANATION OF THE FIELDS

Rank	Site Country / Year	Computer / Processors Manufacturer	Computer Family Model	Inst. type Installation Area	R_{max} R_{peak}	N_{max} n_{half}
1	<u>Earth Simulator Center</u> Japan/2002	Earth-Simulator / 5120 NEC	NEC Vector SX6	Research	35860 40960	1.0752e+06 266240
2	<u>Lawrence Livermore National Laboratory</u> United States/2004	<i>Thunder</i> Intel Itanium2 Tiger4 1.4GHz - Quadrics / 4096 California Digital Corporation	NOW - Intel Itanium Itanium2 Tiger4 Cluster - Quadrics	Research	19940 22938	975000 110000
3	<u>Los Alamos National Laboratory</u> United States/2002	ASCI Q - AlphaServer SC45, 1.25 GHz / 8192 HP	HP AlphaServer Alpha-Server-Cluster	Research	13880 20480	633000 225000
4	<u>IBM - Thomas Watson Research Center</u> United States/2004	BlueGene/L DD1 Prototype (0.5GHz PowerPC 440 w/Custom) / 8192 IBM/ LLNL	IBM BlueGene/L BlueGene/L	Research	11680 16384	331775
5	<u>NCSA</u> United States/2003	<i>Tungsten</i> PowerEdge 1750, P4 Xeon 3.06 GHz, Myrinet / 2500 Dell	Dell Cluster PowerEdge 1750, Myrinet	Academic	9819 15300	630000
6	<u>ECMWF</u> United Kingdom/2004	eServer pSeries 690 (1.9 GHz Power4+) / 2112 IBM	IBM SP SP Power4+, Federation	Research Weather and Climate Research	8955 16051	350000
7	<u>Institute of Physical and Chemical Res. (RIKEN)</u> Japan/2004	RIKEN Super Combined Cluster / 2048 Fujitsu	Fujitsu Cluster Fujitsu Cluster	Research	8728 12534	474200 120000
8	<u>IBM - Thomas Watson Research Center</u> United States/2004	BlueGene/L DD2 Prototype (0.7 GHz PowerPC 440) / 4096 IBM/ LLNL	IBM BlueGene/L BlueGene/L	Research	8655 11469	294911
9	<u>Pacific Northwest National Laboratory</u> United States/2003	<i>Mpp2</i> Integrity rx2600 Itanium2 1.5 GHz, Quadrics / 1936 HP	HP Cluster Integrity rx2600 Itanium2 Cluster	Research	8633 11616	835000 140000
10	<u>Shanghai Supercomputer Center</u> China/2004	Dawning 4000A, Opteron 2.2 GHz, Myrinet / 2560 Dawning	NOW - AMD NOW Cluster - AMD - Myrinet	Research	8061 11264	728400 180000

Peak

40 TF

23 TF

20 TF

16 TF

15 TF

16 TF

13 TF

11 TF

12 TF

11 TF

Linpack

36 TF

20 TF

14 TF

12 TF

10 TF

9 TF

9 TF

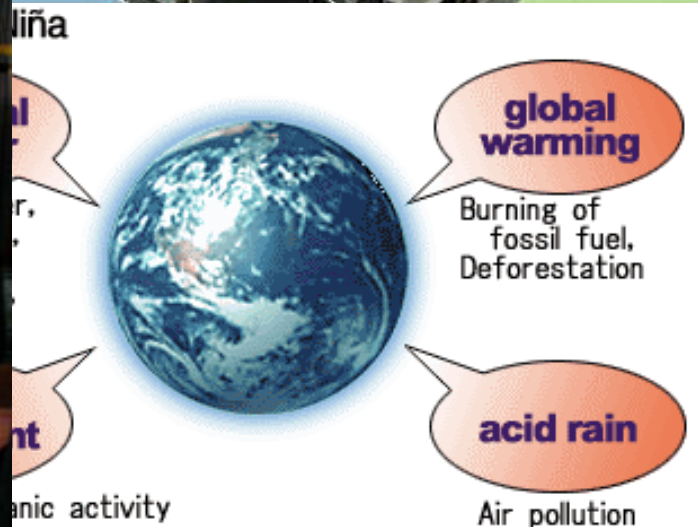

9 TF

9 TF

8 TF

Earth Simulator

- 40 TFlops Peak
- Homogeneous,
Centralized,



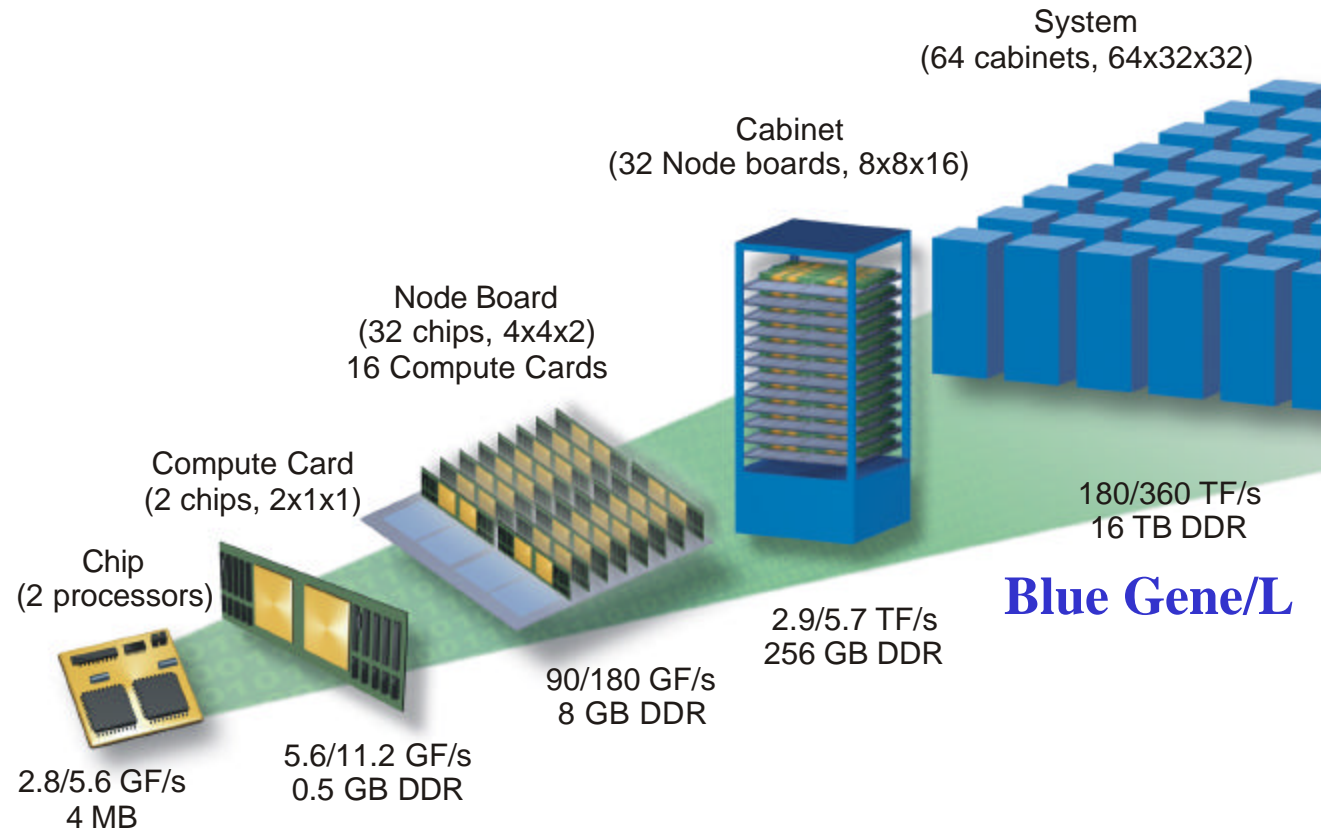
ASCI Purple and Blue Gene/L

■ ASCI Purple (2004)

- ❑ LLNL
- ❑ IBM Power5
- ❑ 12K Processors
- ❑ 100 TFlops

■ Blue Gene/L (2005)

- ❑ LLNL
- ❑ PowerPC
- ❑ 128K Processors
- ❑ Linux
- ❑ 360 Tflops



■ LLNL *Thunder* (2004) 23 TF; 1K 4 IT2, Quadrics

■ Sandia/PSC Cray Red Storm (2004) 40 TF, Opteron, Hypertransport

■ LANL *Lightning* (2004) 11 TF; Linux Networx; 2816 Opteron; Myrinet


■ ORNL Cray X1 (2005) 50 TF; Vector PEs; 2006: 100 TF

■ Dawning (China, 2004) 11 TF; 2560 Operon; Myrinet



CCR Research & Projects


- Ground Water Modeling
- Computational Fluid Dynamics
- Molecular Structure Determination via *Shake-and-Bake*
- Protein Folding
- Digital Signal Processing
- Grid Computing
- Computational Chemistry
- Bioinformatics
- Real-time Simulations and Urban Visualization
- Accident Reconstruction
- Risk Mitigation (GIS)
- Medical Visualization
- High School Workshops
- Virtual Reality




Grid Computing



Grid Computing Overview



Data Acquisition




Advanced Visualization



Analysis




Imaging Instruments



LHC



Computational Resources



Large-Scale Databases

- Coordinate Computing Resources, People, Instruments in Dynamic Geographically-Distributed Multi-Institutional Environment
- Treat Computing Resources like Commodities
 - ❑ Compute cycles, data storage, instruments
 - ❑ Human communication environments
- No Central Control; No Trust


Factors Enabling the Grid

- **Internet is Infrastructure**
 - ❑ Increased network bandwidth and advanced services
- **Advances in Storage Capacity**
 - ❑ Terabyte costs less than \$5,000
- **Internet-Aware Instruments**
- **Increased Availability of Compute Resources**
 - ❑ Clusters, supercomputers, storage, visualization devices
- **Advances in Application Concepts**
 - ❑ Computational science: simulation and modeling
 - ❑ Collaborative environments ® large and varied teams
- **Grids Today**
 - ❑ Moving towards production; Focus on middleware


NSF Extensible TeraGrid Facility

Caltech: Data collection analysis

0.4 TF IA-64
IA32 Datawulf
80 TB Storage




LEGEND

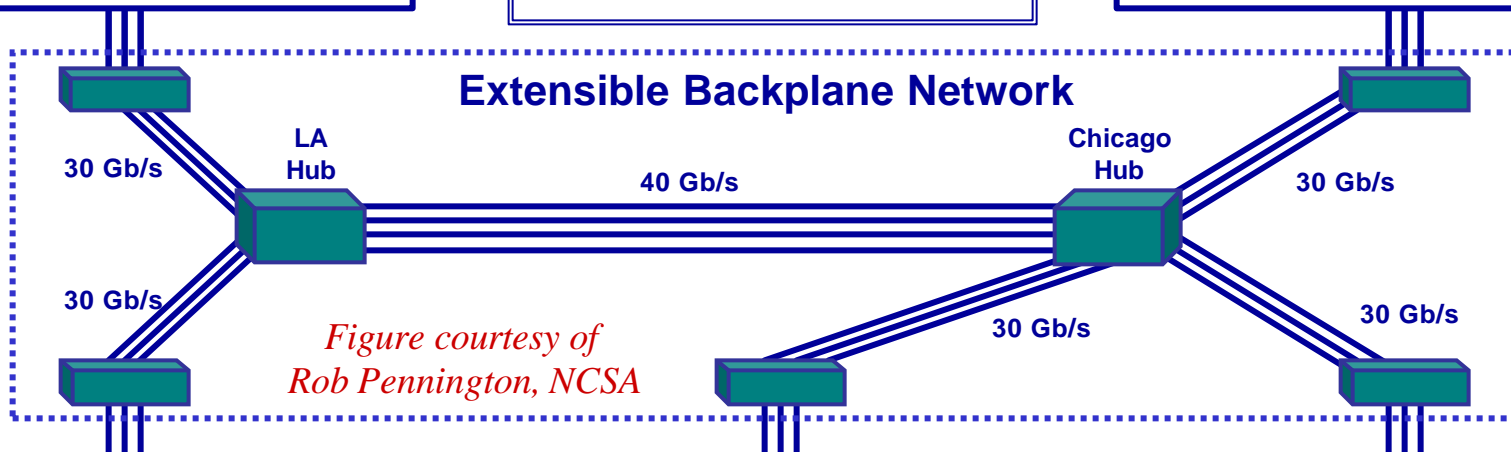


ANL: Visualization

1.25 TF IA-64
96 Viz nodes
20 TB Storage




Extensible Backplane Network




*Figure courtesy of
Rob Pennington, NCSA*

4 TF IA-64
DB2, Oracle Servers
500 TB Disk Storage
6 PB Tape Storage
1.1 TF Power4




SDSC: Data Intensive

10 TF IA-64
128 large memory nodes
230 TB Disk Storage
GPFS and data mining



NCSA: Compute Intensive

6 TF EV68
71 TB Storage
0.3 TF EV7 shared-memory
150 TB Storage Server



PSC: Compute Intensive

Advanced Computational Data Center

ACDC: Grid Overview

Joplin: Compute Cluster

300 Dual Processor
2.4 GHz Intel Xeon
RedHat Linux 7.3
38.7 TB Scratch Space



Nash: Compute Cluster



75 Dual Processor
1 GHz Pentium III
RedHat Linux 7.3
1.8 TB Scratch Space

Mama: Compute Cluster

9 Dual Processor
1 GHz Pentium III
RedHat Linux 7.3
315 GB Scratch Space



ACDC: Grid Portal

4 Processor Dell 6650
1.6 GHz Intel Xeon
RedHat Linux 9.0
66 GB Scratch Space



Young: Compute Cluster

16 Dual Sun Blades
47 Sun Ultra5
Solaris 8
770 GB Scratch Space



Crosby: Compute Cluster

SGI Origin 3800
64 - 400 MHz IP35
IRIX 6.5.14m
360 GB Scratch Space



Expanding

RedHat, IRIX, Solaris,
WINNT, etc

Fogerty: Condor Flock Master

1 Dual Processor
250 MHz IP30
IRIX 6.5



CCR

19 IRIX, RedHat, &
WINNT Processors

Computer Science & Engineering
25 Single Processor Sun Ultra5s

School of Dental Medicine
9 Single Processor Dell P4 Desktops

Hauptman-Woodward Institute
13 Various SGI IRIX Processors

T1 Connection

Note: Network connections are 100 Mbps unless otherwise noted.



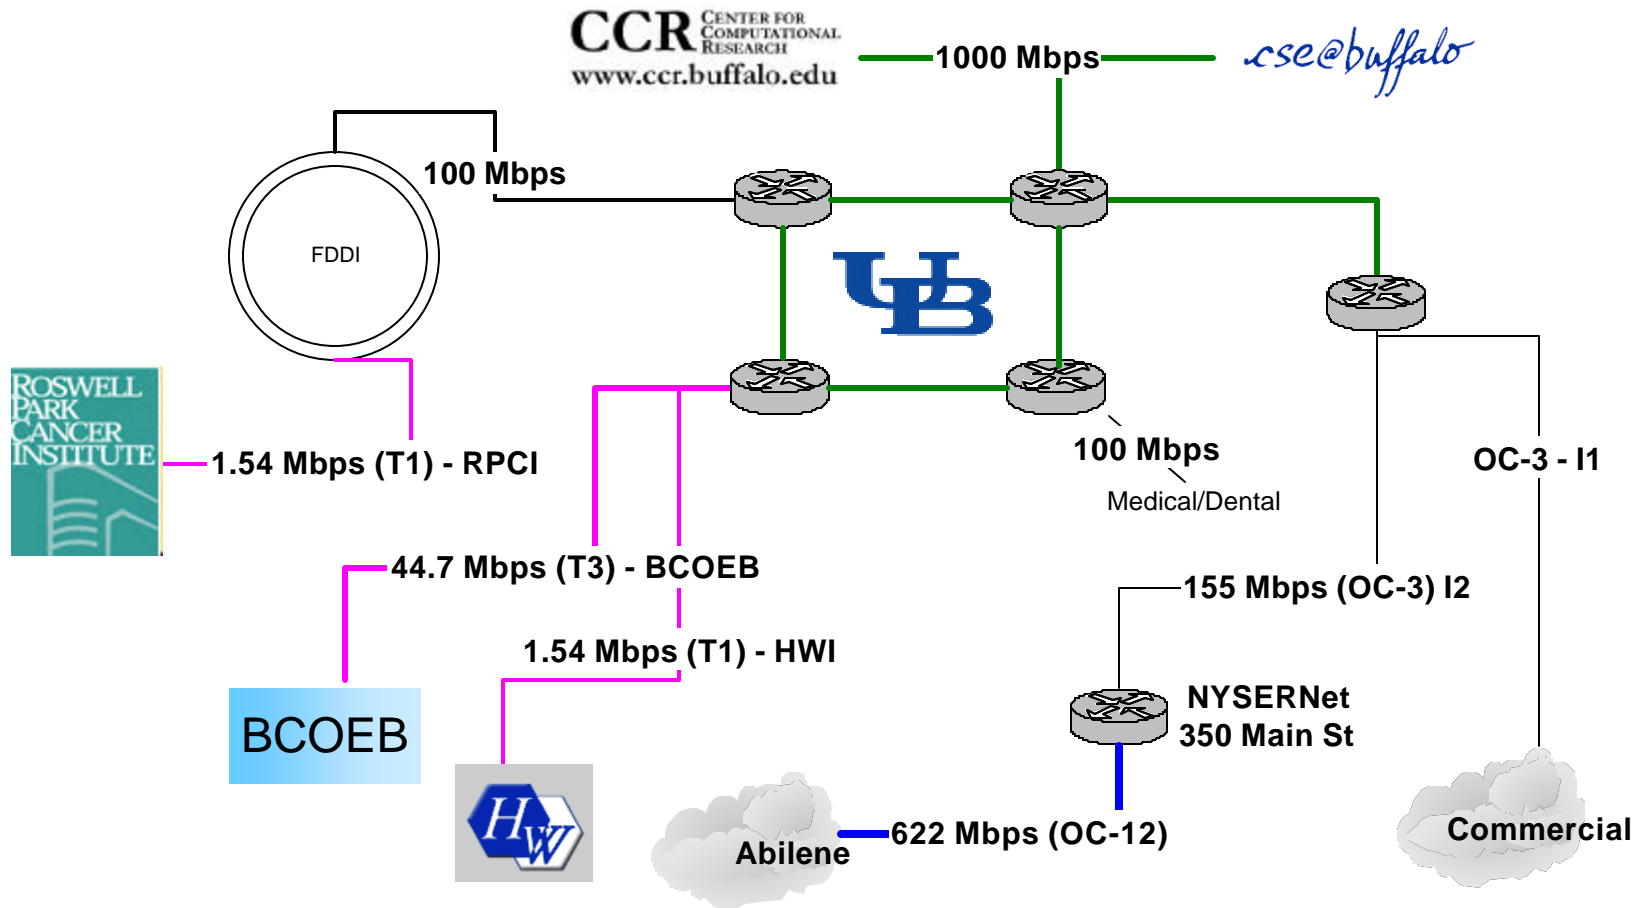
University at Buffalo

The State University of New York

Center for Computational Research

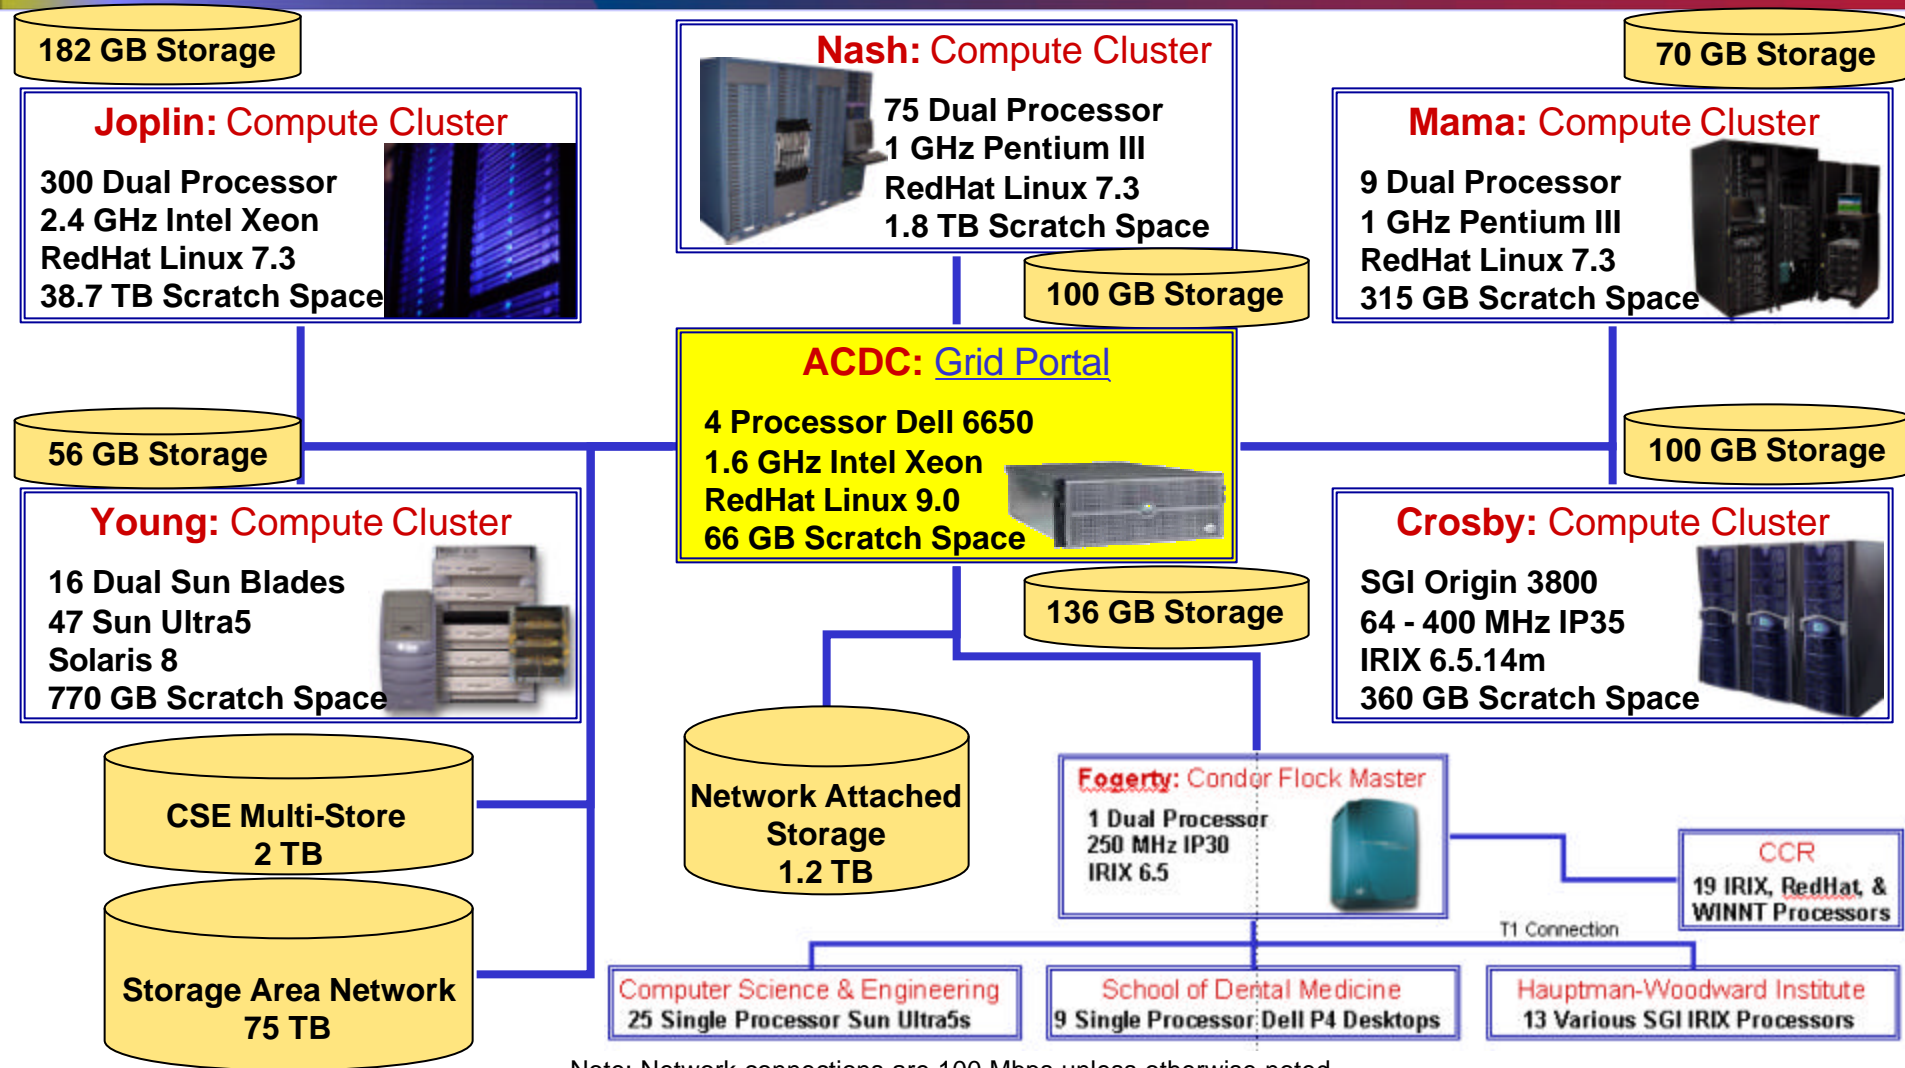
CCR

Network Connections



ACDC Data Grid Overview

(Grid-Available Data Repositories)



Note: Network connections are 100 Mbps unless otherwise noted.



University at Buffalo

The State University of New York

Center for Computational Research

CCR

ACDC-Grid

CCR Grid Computing Services - Microsoft Internet Explorer

CCR University at Buffalo The State University of New York

Center for Computational Research GRID PORTAL

High Performance Grid Computing

Welcome to Grid Computing Services

University at Buffalo Center for Computational Research is currently forming the first Western New York computational grid. The computational grid consist of many supercomputers located at the Center and several other networked supercomputers throughout the Western New York region. These resources will be shared by many researchers from several departments working on a diverse suite of problems including Bioinformatics, Computational Chemistry, and Medical Imaging to name a few.

We also provide grid computing support for the University's Center for Computational Research learning, teaching and research activities plus the infrastructure for both high performance computing and grid enabled software.

Get your "Grid Computing Guide"?

PORTAL LOGIN

- Grid General Info
- Manage Account
- Grid General Info
- Projects
- Resources
- Computational Grid
- Job Submission
- Job Queue Status
- Data Grid
- Network Status
- Running/Queued Jobs
- PBS Job History
- Grid Portal Statistics
- Grid Portal Statistics
- User Info
- Education/Outreach
- Staff Only
- CCR HOME

CCR Grid Computing Services Data Management - Microsoft Internet Explorer

CCR University at Buffalo The State University of New York

Center for Computational Research GRID PORTAL

High Performance Grid Computing

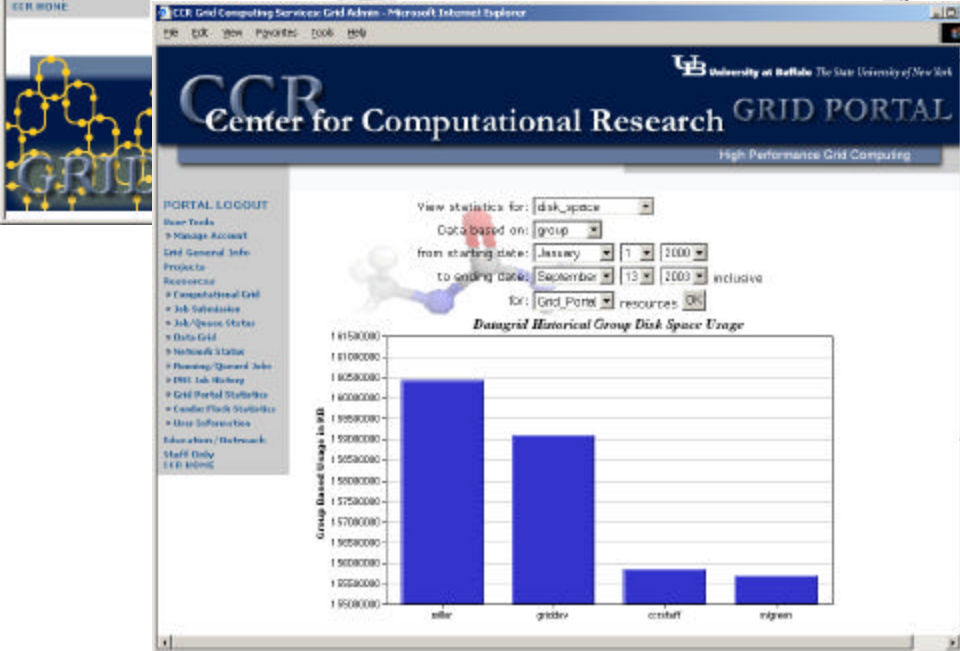
PORTAL LOGOUT

VIEW Group GROUP miller UserList rappleys

reppleye

- KeyMaster
- Morpheus
 - Tank
 - Agent
 - Rabbit
 - Tank
 - Morpheus
 - Oracle.m
 - Nao

Browser view of "miller" group files published by user



CCR Grid Computing Services Grid Admin - Microsoft Internet Explorer

CCR University at Buffalo The State University of New York

Center for Computational Research GRID PORTAL

High Performance Grid Computing

PORTAL LOGOUT

View statistics for: disk_space

Data based on: user

from starting date: January 1 2000

to ending date: September 13 2003 inclusive

for: Grid Portal resources OK

File num	File ID	Filename	Dir ID	Resource ID	Owner	Groupname	Type
1	56033	Cypher.txt	52831	10	mlgreen	griddev	txt
2	56034	Cypher.sh	52858	10	mlgreen	griddev	sh
3	56035	Oracle.asc	52958	10	mlgreen	griddev	asc
4	56036	Cypher.sh	52634	10	mlgreen	miller	sh
5	56037	Rabbit.dat	52830	10	mlgreen	ccrstaff	dat
6	56038	Agent.exe	53064	10	mlgreen	griddev	exe
7	56039	Dozer.sh	52852	10	mlgreen	griddev	sh
8	56040	Nao.asc	52187	10	mlgreen	mlgreen	asc
9	56041	Agent.mpg	52833	10	mlgreen	mlgreen	mpg
10	56042	Tank.txt	52188	10	mlgreen	mlgreen	txt
11	56043	Smith.xls	52258	10	mlgreen	ccrstaff	xls
12	56044	KeyMaster.csh	52186	10	mlgreen	miller	csh
13	56045	Oracle.csh	52632	10	mlgreen	griddev	csh
14	56046	Dozer.xls	52808	10	mlgreen	mlgreen	xls
15	56047	Cypher.exe	52204	10	mlgreen	griddev	exe
16	56048	Rabbit.ppt	52861	10	mlgreen	miller	ppt
17	56049	Nao.dat	52217	10	mlgreen	ccrstaff	dat
18	56050	Cypher.asc	53086	10	mlgreen	griddev	asc



ACDC-Grid Administration

The collage displays four different pages from the CCR Grid Portal:

- Grid Site Administration:** Shows a sidebar with 'PORTAL LOGOUT' and 'User Tools' (Manage Account, Grid General Info, Projects, Resources, Computational Grid, Job Submission, Job/Queue Status, Data Grid, Data Grid Statistics, Network Status, Running/Queued Jobs, PBS Job History, Grid Portal Statistics, Center Fleck Statistics, User Information, Education/Outreach, Staff Only, CCR HOME). The main content area lists 'Users', 'Groups', 'Portal Event Log', and 'Database Job List'. It also includes sections for 'Organizations (add, edit, delete)', 'Resources (view, refresh, ping, delete, create host certificate)', 'Globus Administration', and 'Reports (machine usage, user access to machines, etc.)'.
- Generate Globus grid-mapfile:** A form for generating a grid-mapfile. It includes a text area for 'Optional include file' (set to /home/griddev/www/grid-mapfile.txt) and a text area for 'Optional grid-mapfile path'. A checkbox 'Do not stage this file to the grid nodes' is present. Buttons for 'Generate' and 'Cancel' are at the bottom.
- Create New Database Job:** A form for creating a new database job. It includes fields for 'Job Name', 'Full Path To Script', 'Accepts Arguments' (set to No), 'Run Script' (set to No), and 'Run As User' (set to ccdcc (Admin, Admin)). Buttons for 'Create Job' and 'Cancel' are at the bottom.
- MDS Resource Update Status:** A table showing the status of MDS resources. The current time is 16-September-2003 10:58:12. The table has columns for Resource, Last Updated, Next Update, and Status.

Resource	Last Updated	Next Update	Status
ccrby.ccr.buffalo.edu	16-September-2003 09:15:30	2 minutes	OK
rogerty.ccr.buffalo.edu	16-September-2003 10:45:30	2 minutes	OK
joplin.ccr.buffalo.edu	16-September-2003 10:45:15	2 minutes	OK
mama.ccr.buffalo.edu	16-September-2003 10:45:15	2 minutes	OK
nash.ccr.buffalo.edu	16-September-2003 10:45:15	2 minutes	OK
newus.hwi.buffalo.edu	16-September-2003 10:45:20	2 minutes	OK
yardbirds.ccr.buffalo.edu	16-September-2003 10:45:13	2 minutes	OK
young.ccr.buffalo.edu	16-September-2003 10:45:27	2 minutes	OK



Return to the Grid Resource Admin menu.
Return to the Grid Admin menu.

Advanced
Center for Computational Research
Data
Center



Grid-Enabling Application Templates

- Structural Biology
- Earthquake Engineering
- Pollution Abatement
- Geographic Information Systems & BioHazards



ACDC-Grid

Cyber-Infrastructure

■ Predictive Scheduler

- Define quality of service estimates of job completion, by better estimating job runtimes by profiling users.

■ Data Grid

- Automated Data File Migration based on profiling users.

■ High-performance Grid-enabled Data Repositories

- Develop automated procedures for dynamic data repository creation and deletion.

■ Dynamic Resource Allocation

- Develop automated procedures for dynamic computational resource allocation.

ACDC-Grid Collaborations

- Grid3+ Collaboration / iVDGL Member
- Open Science Grid Founding Participant
 - Monitoring & Information Services, co-chair
 - Security, Tech Working Group Participant
- WNY Grid Initiative
- Grid-Lite
 - HP Labs Collaboration
- Innovative Laboratory Prototype
 - Dell Collaboration
- NE Bio-Grid
 - IBM Research Collaboration
 - MIT, Harvard
- Grid-Based Visualization
 - SGI Collaboration



