PACI and Terascale Update

Robert R. Borchers

PITAC Meeting

May 11, 2001

NPACI

ALLIANCE

PACI Structure

- Leading-edge Site
  - The site with the very large scale computing systems
- Mid-level Resource Sites
  - Partners with alternative or experimental computer architectures,data stores, visualization capabilities, etc and associated training (PACS).
- Applications Technologies
  - Partners involved in development, testing and evaluation of infrastructure from an applications perspective
- Enabling Technologies
  - Partners, generally computer scientists working with computational scientists, developing tools and software infrastructure
- Education, Outreach, and Training partners

Network infrastructure is critical

PACI
Partners for Advanced Computational Infrastructure

An Introduction to PACI

The National Science Foundation's Division of Advanced Computational Infrastructure and Research (ACIR) provides to the national scientific user community support for and access to high-end computing infrastructure and research through its Partnerships for Advanced Computational Infrastructure (PACI) program.

The more than 22 HPC/HTC systems offered through the two national partnerships -- the Alliance and NPACI -- and the recently funded terascale computing system at PSC represent an unprecedented amount of computational resources provided to the American research community by the National Science Foundation. The information contained on this website is your first step to understanding, accessing, and using these resources.

Allocations	Security	Software
User Guides	Computational Resources	Consulting
Training	Science Successes	PACI HotPage

NPACI

ALLIANCE

PACI Allocable Resources

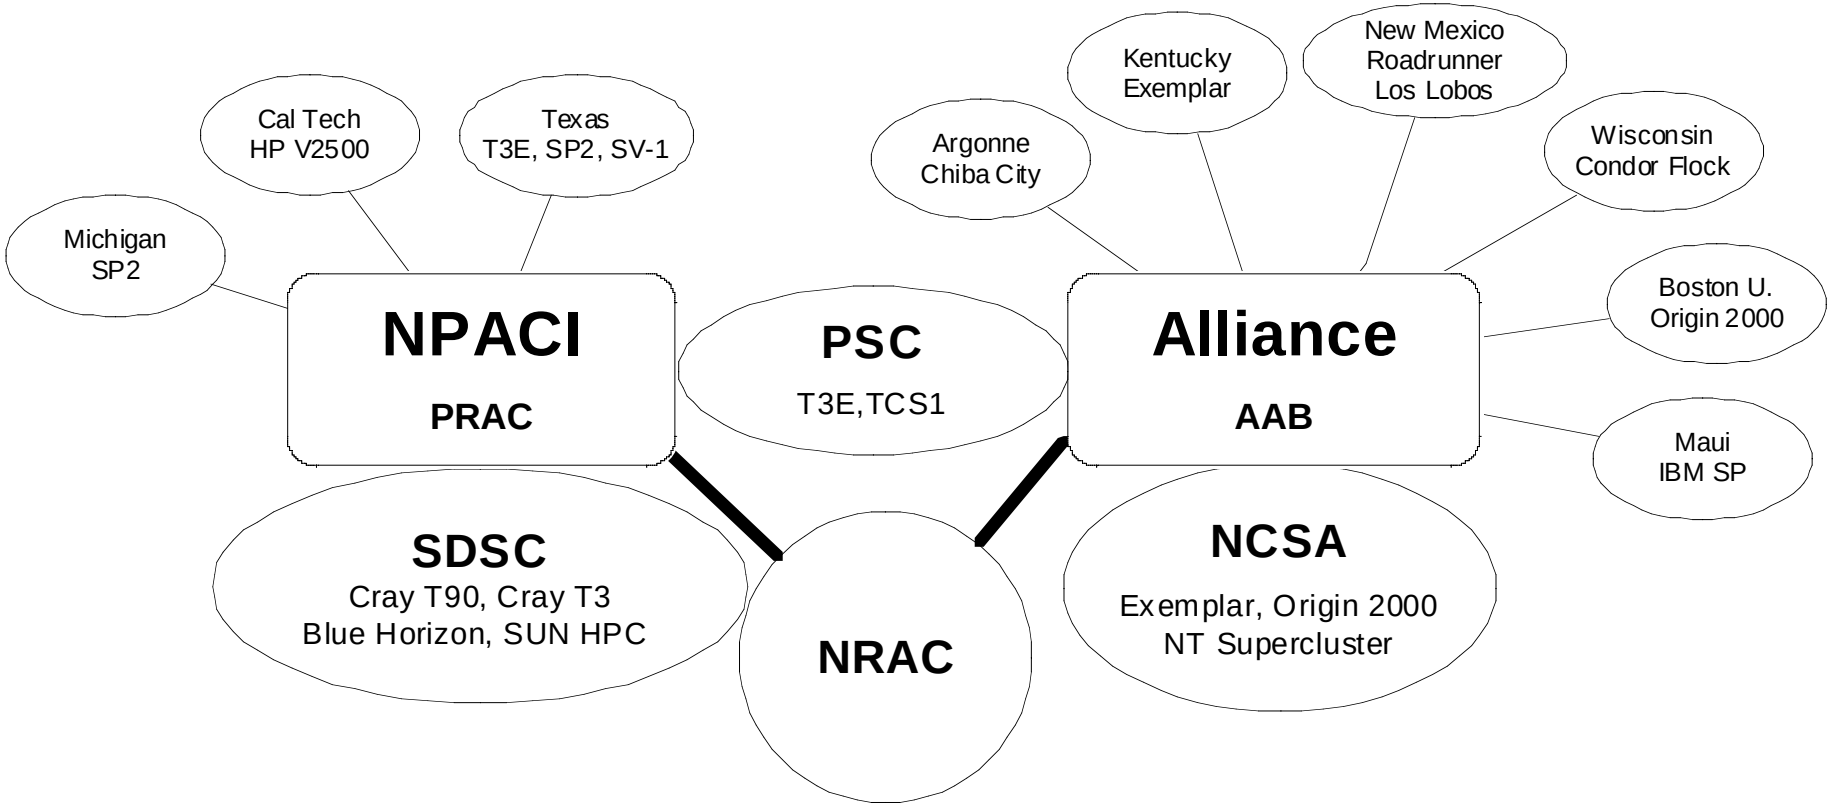

PACI Portal Features

- Grid-enabled through Globus middleware.
- Secure single point login for all PACI resources.
- Continuously updated machine status and job queues.
- User guides for all PACI resources.
- On-line consulting and user trouble tickets.
- Self-paced training materials.
- Documentation on community software codes.
- NRAC electronic proposal submission.
- Success stories across whole PACI Program.

PACI HotPage

home news consulting search feedback

Sunday, March 25, 2001

Login

Services

- Allocations & Accounts
- New User Info
- Documentation
- Consulting
- Training

Resources

- Systems
- Downtimes
- Security
- Applications
- Network Weather Service
- Machine Usage
  - Loads/Queues
  - Processor Node Maps
  - Batch Script Generator

Welcome to the PACI HotPage (v2.1)

The PACI HotPage is designed to *increase the effectiveness of users of PACI's HPC resources*. This page provides links to:

- all NPACI user documentation, including the ability to search (only) the technical documents in NPACI's web.
- news items of current interest
- training and consulting information
- data on computational platforms and software applications
- information about allocations and accounts

In addition, this site provides *active* features such as:

- operational status of computational resources
  - CPU load/percent usage
  - processor node maps
  - queue usage summaries
  - current queue information for all operational platforms
- current MOTDs on all operational platforms
- automated batch script generation for all resources

75%	73%	80%	100%	62%	48%	54%	100%	62%	100%	100%	100%	100%	100%	90%	50%	77%	0%	32%	DOWN	UP	?	Last Update: Mar 25, at 16:30:10 PST
Blue Horizon (horizon)	T3E/272 (golden)	T3E/68 (lonestar)	HPC10K/64 (ultra)	SP/64 (azure)	SP/112 (sp-umich)	T90 (t90)	SV1/16 (aurora)	V2500 (mulder)	V2500 (scully)	NCSA O2K (modi4)	NCSA HP (billie)	BU O2K (lego)	Kentucky N-4000 (nox)	Pittsburgh T3E (jaromir)	NASA O2K (lomax)	NASA O2K (evelyn)	NASA O2K (rogallo)	NASA O2K (sharp)	SDSC (hpss)	UMich (adem)	TACC (dmf)	

TCS-1 Update

- Award Announced 8/3/00.
- 64 node, 256 processor system installed Oct, 2000. ES40 boards with EV68 processors.
- LINPACK Benchmark 263.6 Gflops, (77% of peak) ranked 70 on Top 500 List.
- Passed Acceptance Dec. 22, 2000
- First Advisory Committee Meeting on 2/7/01.
- Friendly User Period through April 1.
- March NRAC awards 2.57M SUs out of 3.79M requested on 64 node system.

TCS-1 Timeline

- 750 4-processor SMP nodes with 4 Gbytes of memory, ES45 boards, EV68 processors, 1GHz
- Nodes start shipping in volume 1st week of June
- Average of 70 nodes/week.
- 256 nodes by 1st week in July; Quadrics Federated Switch tested by then.
- All nodes delivered by Sept. 1.
- 6 Teraflops
- Projected completion by Oct. 2001

NCSA to Build Two New IBM Linux Clusters

- **IA-32 Cluster, 1 Teraflop, under construction**
  - 512 IBM eServer x330 Netfinity thin servers, each with two 1 GHz Intel Pentium III processors and 1 Gbyte memory.
  - Red Hat Linux
  - Myricom's Myrinet 2000 switch, gigabit Ethernet Interconnect
  - 5 Terabytes of Fibre Channel RAID storage
- **IA-64 Itanium Cluster, 1 Teraflop, this summer**
  - 160 IBM X-Series 800 MHz IA-64 dual processor systems with 2 GB of memory

NPACI "Rocks" Cluster Toolkit

- Phil Papadopoulos, group leader for Distributed Computing at SDSC
- Only a handful of simple steps are required to bring up a full-featured cluster
- Rocks installs the Maui Scheduler, Portable Batch System, and other tools in a complete production environment
- Upgraded to RedHat 7.0
- Supports multiple high-performance interconnects such as Myrinet (today) and Servernet II (soon)

Data Intensive Computational Environments

SDSC
SRB
Storage Resource Broker

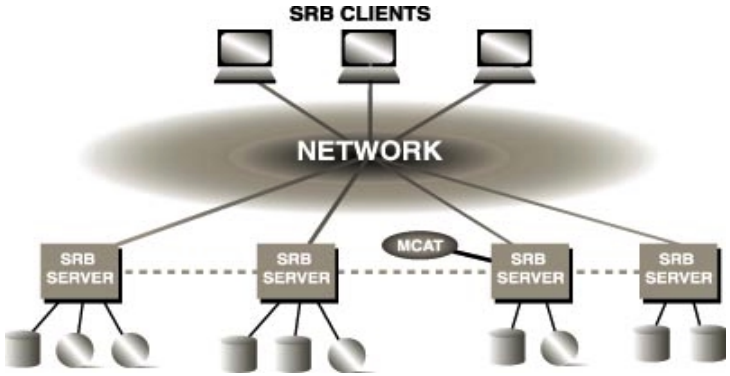

- **Client-server middleware**
- **Provides a uniform interface for connecting to heterogeneous data resources over a network.**
- **Provides access to data sets and resources based on their attributes rather than their names or physical locations**

SRB Data Collections

Collection Site	Project	Archive (TB)	Database (GB)
Caltech	Digital Sky images / NPACI-DICE	2-20	2
SDSC	CEED / ESA	1	2
SDSC	PDB	0.5	2
SDSC	NARA – USPTO patents	0.3	70
SDSC	Human Brain Project / NPACI-Alpha NS	1	10
SDSC	Molecular Structures / NPACI-Alpha MS	1	10
SDSC	Visualization image collection	0.5	5
SDSC	SRB Production system / NPACI-DICE		75
UCB	Elib flora collection / NPACI-DICE	1	60
UCLA	Human Brain Project / NPACI-NS	1	2
UCOP	Art Museum Image Consortium / CDL	1.5	30
UCSB	Alexander Digital Library / NPACI-DICE	2	2
UCSC	REINAS / NPACI-ESS	0.1	1
U Maryland	HPSS federation / NPACI-DICE	1	1
U Michigan	UMDL / NPACI-DICE	0.1	30
U Wisconsin	Digital Insight / NPACI-EOT	10-20	5
Washington U	Human Brain Project / NPACI-NS	1	10
UCSD	PRDLA - proposal	0.5	10
UCOP	CDL backup - proposal	5	10

PACI EOT Launches $2M Teacher Training Program at SC2000

- $1.13M grant from EHR.
- State-of-the-art laptops from Compaq
- 25 Teams of 4 teachers receive one week of immersive training at SC2000
- 18-month program
- Modeling package Stella from High Performance Computing
- Microsoft Excel
- Mathematic from Wolfram Research
- Objective is to design classroom modules available to all high schools across the country

NCSA Receives $300K NSF Planning Grant for NEES-grid

- NEES grid is a national virtual laboratory for earthquake engineering.
- Built on existing grid Globus technologies developed at Argonne and USC
- Will allow seamless sharing of experimental equipment, computational resources and data.
- NEES grid will be a tele-observation and tele-operation tool.

Distributed Terascale Facility

- DTF Description
  - multi-site "distributed facility" connected by ultra high-speed networking
  - one single-site computing system capable of five or more teraflops per second (peak) performance
  - system will be embedded within an overall system that also provides sophisticated data handling and interaction with remote sites
  - facility will include substantial support for accessing, analyzing, processing, transmitting, and visualizing multi-terabyte data collections
- Solicitation NSF 01-51 posted Jan. 19, 2001
- Proposal Deadline Apr. 19, 2001
- Panel met May 3&4, 2001
- Site Visit planned for June 5, 2001