Center for Computational Sciences

ORNL
Center for Computational Sciences

Topics

- Mission
- Supercomputers and plans
- Storage and networks and their use
- New building
- Special mention

CCS: Principal Resource for SciDAC

Resources:

- 4.5 TFlops IBM Power4
- 1 TFlops IBM Power3
- 0.5 TFlops Compaq
- High performance storage
- Visualization hardware
- OC192 Network

Applies dedicated computational resources to accelerate breakthrough science in support Scientific Discovery through Advanced Computing (SciDAC)

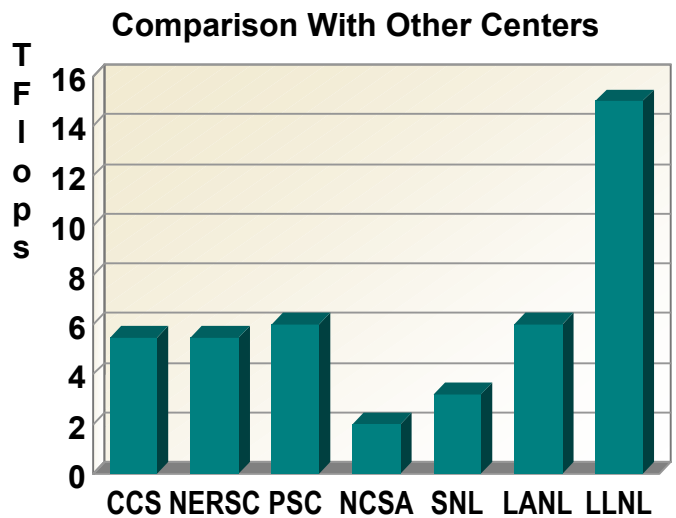

Center for Computational Sciences

CCS a Partner for SciDAC Success in Key Projects

- Scientific Applications
  - **Climate Simulation**
  - **Fusion**
  - **Astrophysics**
  - **Chemistry**
- Collaboratories
  - **Science Grid**
  - **Earth Systems Grid**
- Middleware & Network Research
  - **Network Dynamics**
  - **Net100**
  - **SSFnet**
- Prototype Topical Center
- Computer Science
  - **Scalable Systems Software**
  - **Common Component Architecture**
  - **Performance Science and Engineering**
  - **Scientific Data Management**
- Applied Mathematics
  - **PDE Linear/Nonlinear Solvers and Libraries**
  - **Unstructured Grids**
- Scientific Application Pilots
  - **Astrophysics**
  - **Climate**
  - **Fusion**

Partnerships with 13 laboratories and 50 universities

CCS – Major Computers

Current and Future

- **Falcon Compaq Alphaserver (Procured- FY 1999: Planned Retirement – FY 2004)**
  - **Principal resource for computational biology**
- **Eagle IBM SP (Procured- FY 1999: Planned Retirement – FY 2004)**
  - **Principal resource for nanoscale science**
- **Cheetah IBM Regatta (Procured- FY 2002: Planned Upgrade-FY 2003: Planned Retirement –FY 2006)**
  - **Principal resource for DOE SciDAC**
- **Cray X1(Procured-FY 2003:Planned Upgrade FY2003-2004: Planned Retirement – FY 2007)**
- **Blue Gene/D (Designed- FY2002-4: Planned Procurement- FY 2005; Retirement –FY 2010)**
  - **Specifically designed for nano, biology and climate**

CHEETAH

- 4.5 TFlops, IBM Power4 Regatta System
- 27 computational nodes, 32 processors each = 864
- 5.2 GFlops/sec peak processor speed
- 1.2 TB total memory
- 40 TB of disk space in GPFS
- Upgrade to 10 TFlops in 2003

- Upgrading to Federation switch this summer
  - Beta test; some disruption possible
  - Much faster and with much reduced latency (2 GB/sec vs 180 MB/s, 4 paths vs 2, 9 microsec latency vs 25)

Cray X1

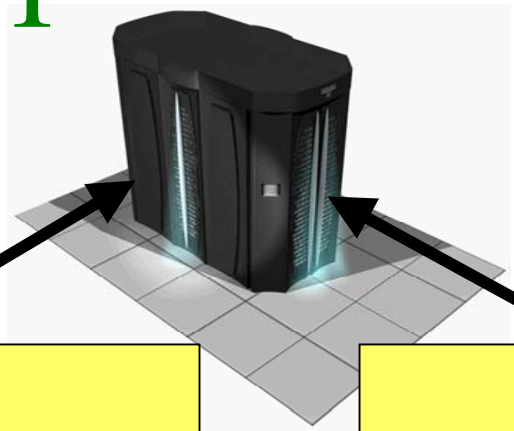

Cray PVP

- Powerful vector processors
- Very high memory bandwidth
- Non-unit stride computation
- Special ISA features

- Modernized vector ISA

T3E

- Extreme scalability
- Optimized communication
- Distributed memory
- Synchronization features

- Improved via vectors

Extreme scalability with high performance vector processors

Phase 1 – First Quarter 2003

- 32 Vector Processors
  - 8 nodes, each with 4 processors
- 128 GB shared memory
- 8 TB of disk space

400 GigaFLOP/s

Phase 2 – Summer 2003

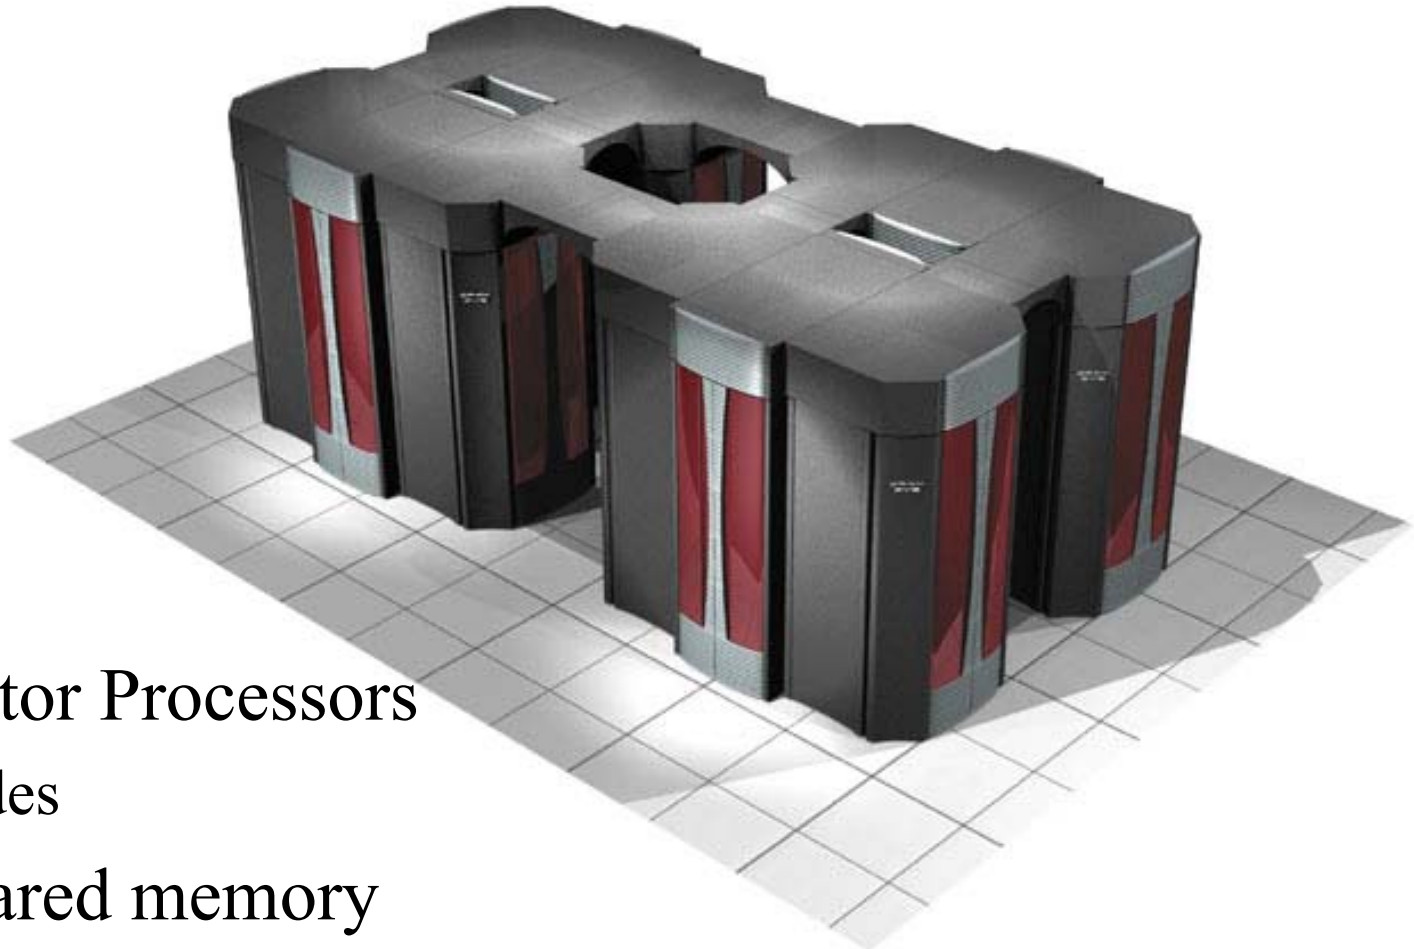

- 256 Vector Processors
  - 64 nodes
- 1 TB shared memory
- 20 TB of disk space

3.2 TeraFLOP/s

Blue Gene/D

- **Balanced and tunable ultra-scale parallel computer specifically designed and developed for:**
  - **Nanoscale Science**
  - **Genomes to Life**
  - **Climate Change initiatives**
- **Initial contract signed with IBM to jointly develop Blue Gene/D specifically for GTL, Nano and Climate**
  - **Five ORNL-IBM Fellowships established to attract the best**
- **Goal at least 100X performance improvement over NERSC-3 within approximately the same cost;**
- **Planned FY 2004 procurement with initial delivery in FY 2005**

Storage and Networks and Their Use

- Mass Storage
  - Currently 2 silos in use – 11,000-tape capacity
  - 2 more will be added in June
  - With new-technology drives, capacity > 2 PB
- OC192 available (10 gigabits/second)
  - Connected to Internet 2 and carrying traffic
  - Fully using the bandwidth, especially wide-area, is a challenge
  - Several projects underway to optimize usage
- ESnet link OC12 (620 megabits/second) now; going to OC48 or OC192 this year, funds permitting

Construction began in 2002 and ready for installation in June 2003

Center for Computational Sciences

- Space and power for world class facilities
  - 40,000 ft2 Computer Center
  - Largest in DOE/SC
- Office space for 500 staff members
- Classroom and training areas for users
- High ceiling area for visualization lab (Cave, Power Wall, Access Grid, etc.)
- Separate lab areas for computer science and network research
- Strong university partnerships

Special Mention

- Schedule through September
  - Really soon – Phase 1 Cray; existing computer room
  - Summer – Cheetah upgrade; some disruption possible
  - Summer – move to new building; shut entire center for 1-2 weeks
  - Summer – Phase 2 Cray in new building.
- We must know what you need and when
  - To define, justify and (try to) fund equipment purchases
  - To prototype and harden capabilities, especially in storage and networking
- Need information on:
  - Rate of data generation – simulation and analysis
  - Capacity and duration
  - Bandwidth to storage, other servers and other sites

Questions?

CCS Major Computers

	Intel Paragon XP/S-35	Intel Paragon XP/S-150 MP CCS-1	IBM RS/6000 SP "Eagle" CCS-2	Compaq Alpha "Falcon"	IBM Power4 "Cheetah" CCS-3
Year of Installation	1992	1995	1999	2000	2002
Compute Processors	512	2048	704	265	865
Processor Technology	Commodity CMOS	Commodity CMOS	Commodity CMOS	Commodity CMOS	Commodity CMOS
Peak System Perform.	35 Gflop/s	150 Gflop/s	1,080 Gflop/s	342 Gflop/s	4493 Gflop/s
Architecture	Distributed memory	1024 nodes with 2 processor SMP	176 nodes with 4 processor SMP	64 nodes with 4 processor SMP	27 nodes with 32 processor SMP
System	Fully integrated custom system with commodity CPU and memory	Fully integrated custom system with commodity CPU and memory	Loosely integrated system with commodity system components	Loosely integrated system with commodity system components	Loosely integrated system with commodity system components
System Software	Mach microkernel based OSF/1 Unix	Mach microkernel based OSF/1 Unix	IBM AIX Unix with PSSP parallel system	Compaq Tru64 Unix with RMS parallel system from QSW	IBM AIX Unix with PSSP parallel system
Floor space	96ft^2	203 ft^2	456 ft^2	280 ft^2	1200 ft^2
Power consumption	94 kW	200 kW	135 kW	66 kW	