HEPiX Spring 2005 Overview

Chris Hollowell <*hollowec@bnl.gov*>

- Site Report Highlights
- Monitoring Systems
- Security@SLAC
- Storage Systems
- XROOTD@RAL
- Scientific Linux Status Report
- Fermilab Opteron Evaluation
- Batch Systems Workshop

Site Report Highlights

- **GridKa@FZK**

  -Using ceiling power rails

  -Cooling issues: solved with Knurr water cooled cabinets

  -Scalability problems with OpenPBS: now using PBSPro

  -3 GPFS clusters consisting of 12 systems each: ~300 TB total

  -IBM 3584 tape storage (LTO-1 and LTO-2 tapes/drives)

    -Tape quality differs between vendors

Site Report Highlights (Cont.)

- **TRIUMF**

  -Budget decrease

  -Some bad experiences with SATA reliability: believe these disks require more quality control, screening and repair mechanisms

  -Recommend RAID6 for large arrays

- **CASPUR**

  -Using Sun Grid Engine

  -8 TB of storage available via IBM SANFS: in semi-production

  -Using their own Linux distribution: BigBox

Site Report Highlights (Cont.)

- **SLAC**

  -6 month shutdown due to power accident: now there is quite a large emphasis on safety

  -Facing 5% staff reduction due to power cost increase and DOE budget reduction

  -Lab is beginning to diversify beyond HEP

  -PetaCache Project:

    -Store data in a farm of memory (RAM) servers for extremely fast retrieval (100 microsecond latency): 1 TB prototype cluster in place

    -Demonstrated scaling to 100 servers

    -Data managed by XROOTD on x86_64 Solaris

Site Reports Highlights (Cont.)

- **Fermilab**

  -Finances tight: 5% staff reduction, BTeV canceled

  -dCache in use at USCMS Tier1 facility: implemented resilient dCache configuration using compute node local data storage

- **CERN**

  -Migrating Linux systems from RH7.3 to SLC3

  -AFS authentication migrated to K5

Site Report Highlights (Cont.)

- -Random AFS client crashes seen: determined to be an incompatibility between new servers and older clients
- -Noted a 50% decrease in spam during the first half of 2005
- -No longer purchasing tower systems
- -Many compromised accounts on lxplus systems
  - -In one instance intruder used an insecure setuid LSF application to gain root
  - -Now blocking P2P and IRC traffic at their firewall

Site Report Highlights (Cont.)

- **RAL**

  -RH7.3 installed on NFS servers: SL3 NFS server performance poor. Looking into migrating to SL4

- **GSI**

  -Evaluating the replacement of NIS with LDAP

  -Successfully tested GEANT3 and ROOT under 64-bit Debian on Opteron

- **INFN**

  -Torque+Mauii not scalable: migrating to LSF 6.0

Monitoring Systems

- **CERN LEMON (LHC Era Monitoring)**

  -Developed at CERN

  http://cern.ch/lemon

  -C++ API available for writing one's own monitoring sensors

  -Total of 260 sensors currently available in the system (monitoring of system utilization, hardware and software problems, system security, etc.)

  -Allows for reboot scheduling based on metrics: useful for kernel upgrades.

Monitoring Systems (Cont.)

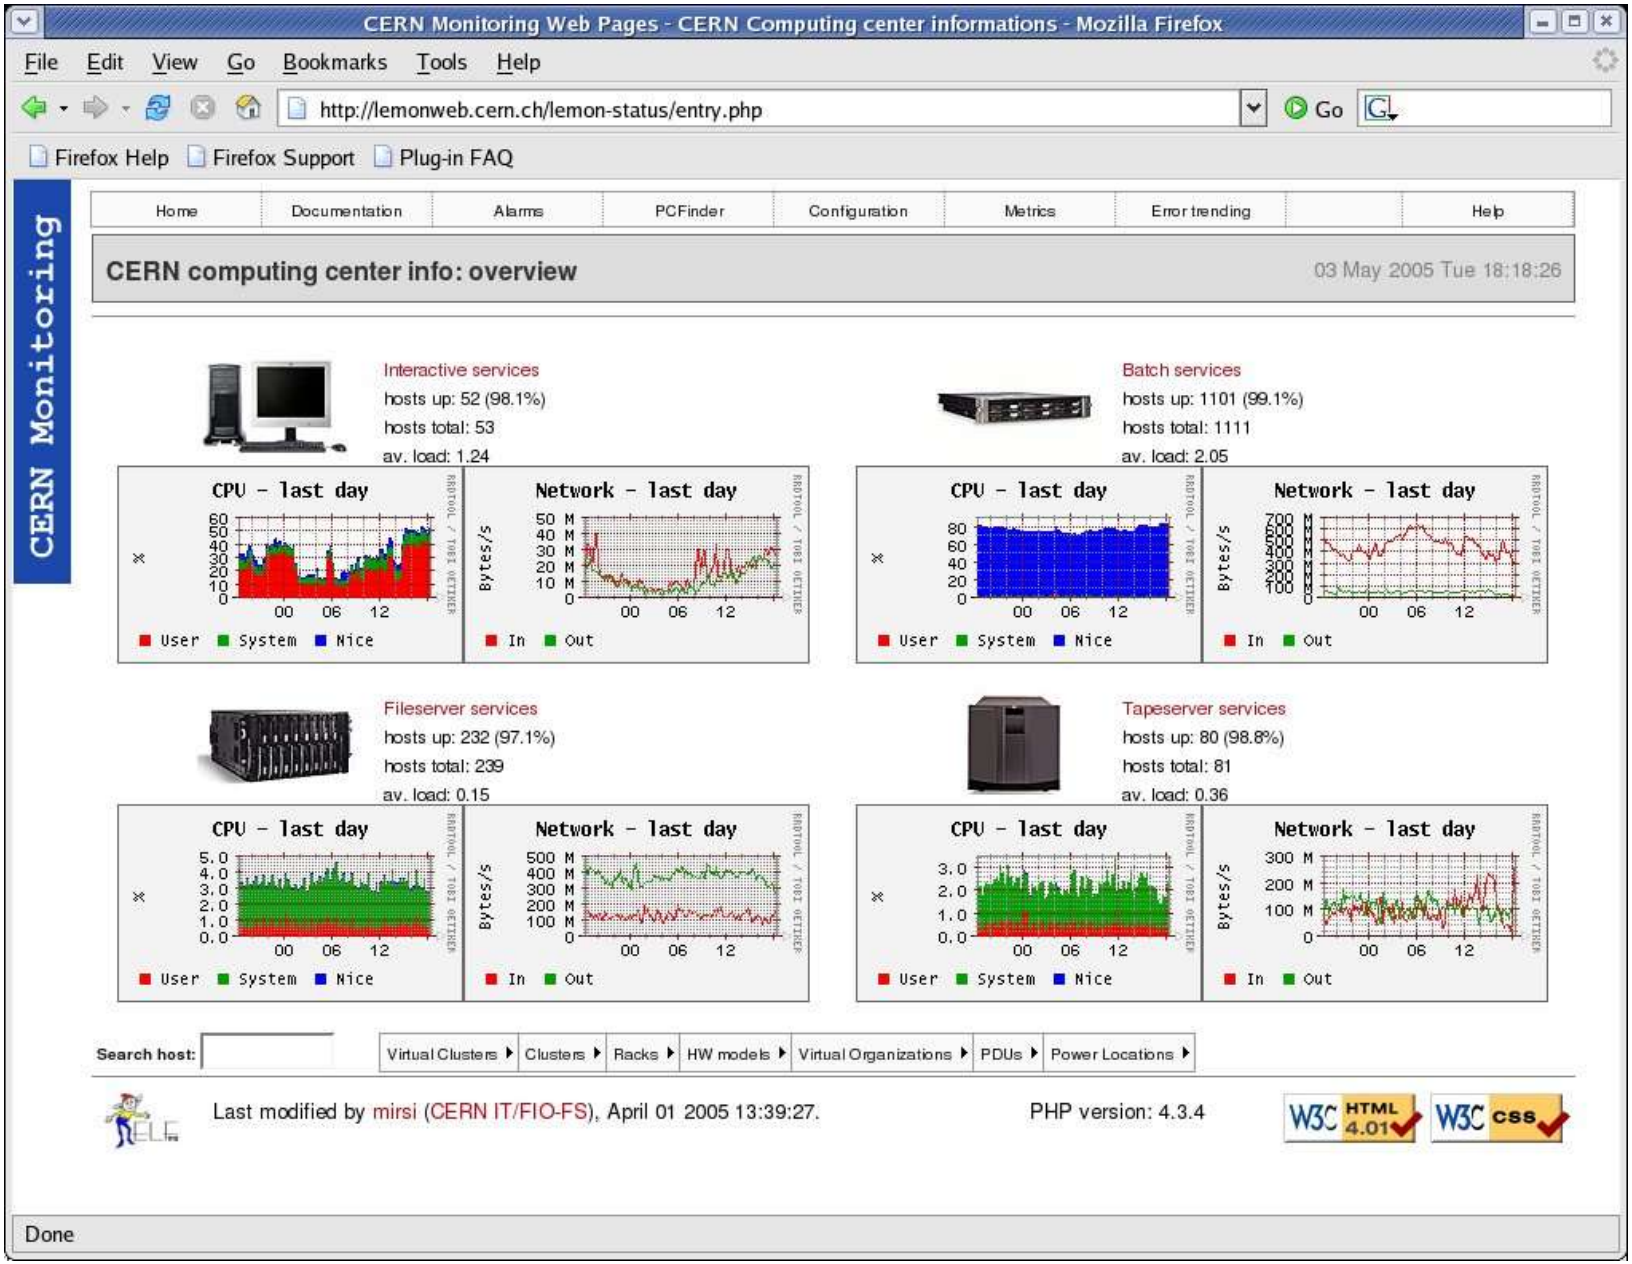

Monitoring Systems (Cont.)

-For defined values of measured metrics, an actuator can be called with a predefined action

-Example - SSH daemon not running: restart

- **SLAC Experience with Nagios**

-Open source monitoring software:

http://www.nagios.org

-Fairly easy to write one's own monitoring extensions (via perl modules)

-Third-party software has the ability store metric history in a MySQL database (PerfParse) or in RRDs (nagiosgraph)

Monitoring Systems (Cont.)

-Scalability issues with 1.x web interface: 2.x (still in beta) performs much better

- NGOp at Fermilab

  -Developed at Fermilab: Next Generation OPerations

  -Scalability claims:

    Ganglia: 500 systems (sic)

    Big Brother: 75 systems

    Nagios: 650 systems

    NGOp: 2000 systems

  -Work needs to be done to make NGOp functional on a grid level

Monitoring Systems (Cont.)

-NGOp metric naming conventions need to be standardized

-Has the ability to make distinctions between major and minor problems (i.e.: one host down vs. all hosts down)

Security@SLAC

- SLAC has witnessed ”session hijacking”
  - -One-time passwords aren't a complete solution
- Number of account compromises has decreased
  - -The group of people responsible seem to have quieted down lately
- On the increase:
  - -Phising
  - -Pharming (DNS Server Poisoning)
  - -Spyware

Security@SLAC (Cont.)

- **New technologies, new risks:**

  -RFID Tags: RFID codes can be changed

  -VoIP: allows for manipulation of caller-ID

  -Firewire devices: firewire specification (IEEE1394) allows devices to read/write directly to memory, bypassing OS limitations. Malicious devices could be used to access sensitive information or compromise a system

Storage Systems

- **CASPUR**

  -Using Panasas and Lustre

  -Out of NFS, Lustre, and Panasas, benchmarks indicate Panasas performs the best

- **Fermilab USCMS Update on Use of IBRIX Fusion FS**

  -IBRIX Fusion FS is a scalable parallel filesystem

Storage Systems (Cont.)

- Decided to purchase IBRIX because:
  - It requires no specialized hardware
  - Provides NFS client access to storage
  - IBRIX was responsive to issues
- IBRIX Fusion FS features:
  - Distributed metadata
  - Automatic server failover
  - Dynamic Load Balancing
- Client systems may use ”IBRIX Lite” thinclient to access data
  - Driver can route requests directly to correct segment servers

Storage Systems (Cont.)

- Experienced many problems with IBRIX Fusion 1.2
  - Fail-over mechanisms didn't function as expected
  - Issues with NFS access
- Due to stability issues, stopped using IBRIX Fusion in August 2004: reassigned associated hardware to standard NFS NAS duty
- Installed version 1.3 in Feb 2005
  - Stability greatly improved
- Planning on making 20 TB available before data taking begins

Storage Systems (Cont.)

-IBRIX Lite package is kernel version dependent: IBRIX has been providing new versions of the package for kernel updates in a timely manner

XROOTD@RAL

- XROOTD: eXtended ROOT Daemon
  - -Developed a SLAC
- Part of ROOT since September 2004 (version 4.01-02)
  - -Also available independently:
    - http://xrootd.slac.stanford.edu
- Goals:
  - -High performance
  - -Fault tolerance

XROOTD@RAL (Cont.)

- **Performance features:**

  -Connection multiplexing: multiple data requests from a particular host channeled through one TCP connection

  -Request redirection/load balancing

- **Consists of data servers and redirector/load balancing servers**

- **CacheFS:**

  -Allows XROOTD to access multiple disks on a host using a single namespace

  -Distributed with XROOTD

XROOTD@RAL (Cont.)

-Provides several scripts to manage namespace

-Namespace served by XROOTD is a tree of links

- **MPS Tools**

-Migration, Purging and Staging Tools

-Used to interface XROOTD with mass storage systems

-Makes calls to user defined commands, effectively allowing interface with any storage system

XROOTD@RAL (Cont.)

-Supports pre-staging and stage request queuing

- Setup at RAL

-Two redirector/load balancing servers (olb managers) implemented via a round-robin DNS alias

-If a server fails, and files from this server are requested, they are automatically staged from mass storage

-Using CacheFS to aggregate local filesystems on hosts

-26 data servers managing ~70 TB of disk

XROOTD@RAL (Cont).

- **Users aren't heavily impacted by disk/server failures:**

  -Data is retrieved from another host or from mass storage: jobs don't fail
  -Makes scheduling downtime for one or more systems easier

Scientific Linux Status Report

- Initial release 1 year old
- List of contributors to SL growing
  - More still needed
- Commitment to keeping system ABI compatible with RHEL
- Monthly SL video conferences
- Planning a Scientific Linux Workshop for Fall 2005
- Attempting to build a RHEL compatibility testing tool

Scientific Linux Status Report (Cont.)

- **File downloads from FNAL Scientific Linux FTP server:**

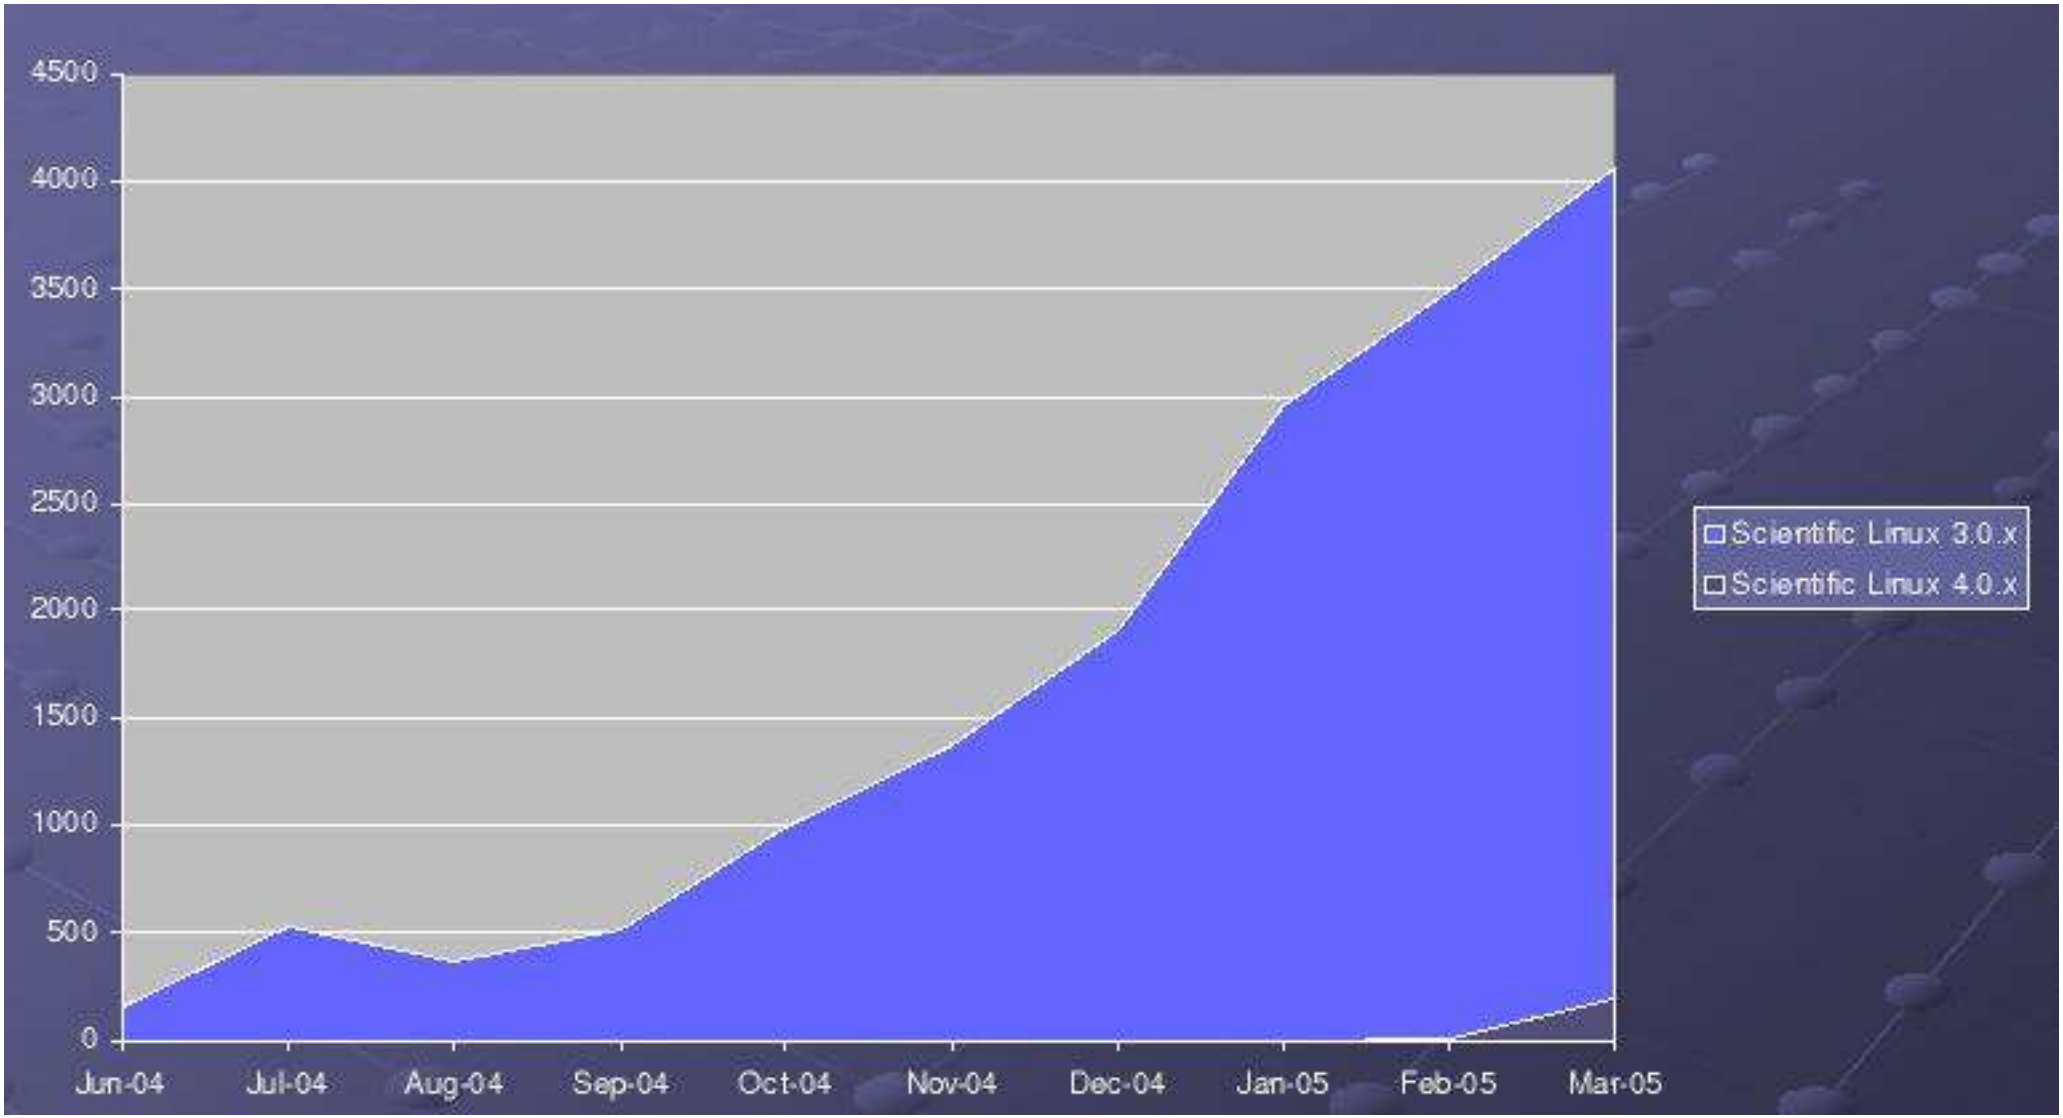

Scientific Linux Status Report (Cont.)

- Scientific Linux CERN
  - -RH 7.3 EOL'd end of 2005 at CERN
  - -SL3/RHEL3/SLC3 compatibility a goal
    - -Only one documented SLC3 incompatibility with SL3: libstdc++ regression for POOL
      - -CERN's ”fault”, but required for certification at CERN
    - -Some overzealous/accidental dependencies on /etc/redhat-release encountered
  - -Concerned over many OS choices for experiments in the future (SL3, SL4, SL5) and the introduction of x86_64
  - -CERN has no plans for ”jumping on SL4”

Scientific Linux Status Report (Cont.)

-Need a supported/new and stable OS for LHC startup in 2007: release schedule extrapolation indicates SLC5 is the likely candidate Devoting resources to creating/certifying SLC4 is likely wasteful

-CERN's claim: if CERN doesn't move to SL4, it will not become the ”default HEP platform”

Fermilab Opteron Evaluation

- Several AMD Opteron and Intel Xeon based systems evaluated.
- Differences in instruction set:
  - -Xeon/EM64T supports SSE3 instructions
  - -Opteron supports 3DNow instructions
- Dell doesn't offer Opteron systems
- Opteron memory architecture more advanced than Xeon's:
  - -Each CPU has its own memory modules

Fermilab Opteron Evaluation (Cont.)

-CPUs access each other's memory via HyperTransport Bus

-In conjunction with an OS with NUMA support, minimizes memory contention between processors

- Linux kernel 2.6.9 required to take full advantage of NUMA with Opterons
- Different kernels required for Xeon and Opteron processors
- Successfully ran almost all of their 32-bit applications under x86_64 SL

Fermilab Opteron Evaluation (Cont.)

- **CMS ROOT benchmark results: 64-bit operation yields gains of ~40% for Opterons**

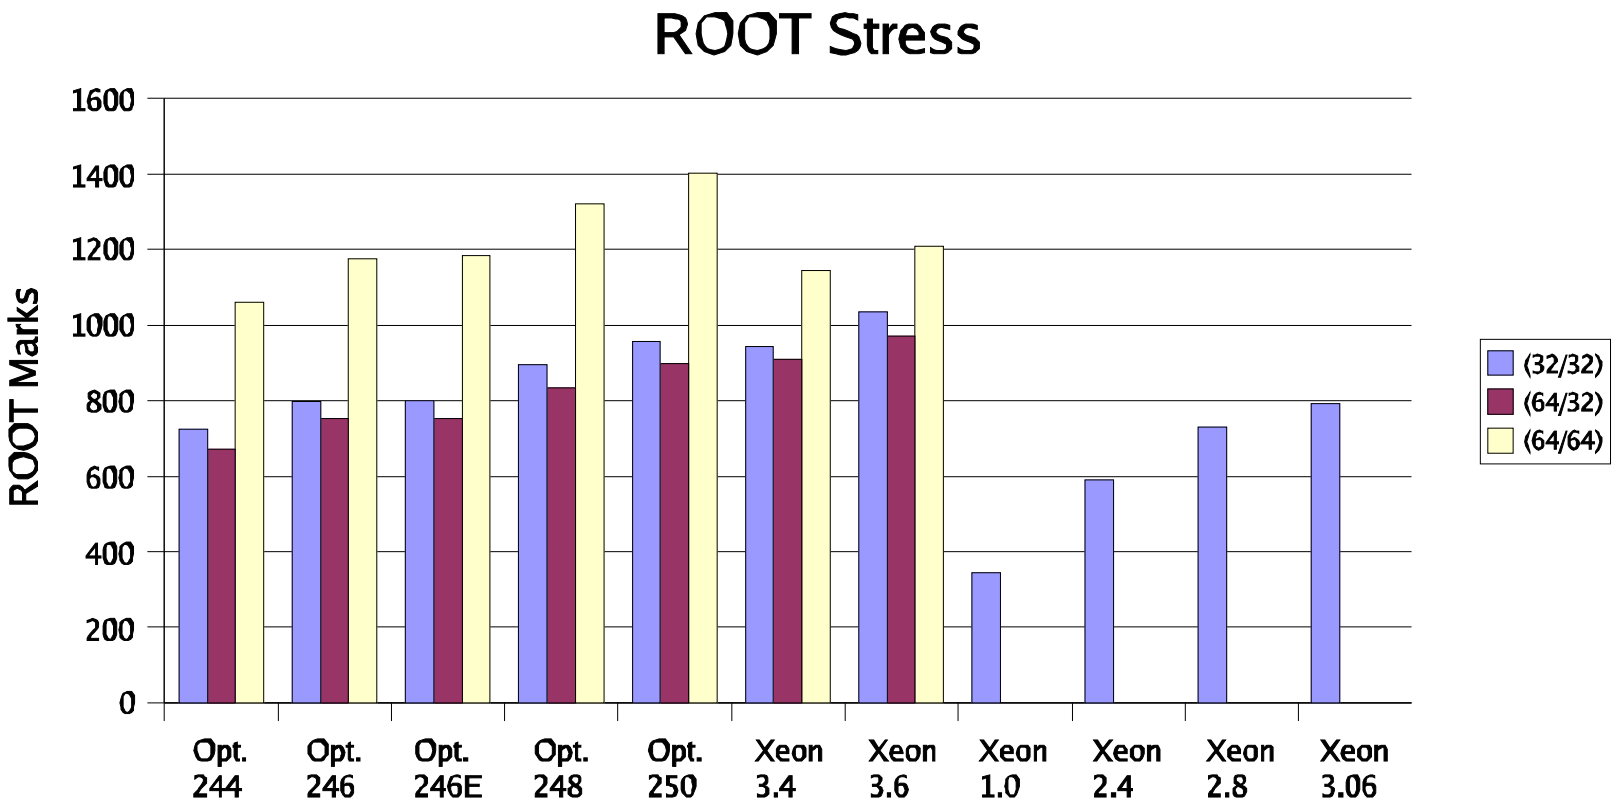

Fermilab Opteron Evaluation (Cont.)

- Version 8.1 of the Intel Fortran and C compilers have the ability to output SS3 instructions for code optimization: can make software built with these compilers incompatible with Opterons
- Power consumption

  -In general Opterons draw 10-27% less power at full load than Xeons.

  -Xeon power consumption varies more with CPU utilization; Opteron power draw relatively static. Xeons draw less power than equivalent Opterons when idle

Fermilab Opteron Evaluation (Cont.)

-Low voltage Opteron 246HE draws 10-15% less power than high voltage Opteron 246

Vendor	Processor	Idle	Loaded
HP	Opt244	1.54	1.74
Rackable	Opt246HE	1.44	1.61
IBM	Opt246	1.55	1.80
Penguin	Opt248	1.70	1.98
ASA	Opt248	1.91	2.20
CSI	Opt248	1.95	2.30
Sun	Opt248	2.00	2.35
Koi	Opt250	2.07	2.44
Dell	Xeon3.6	1.66	2.75
Koi	Xeon3.6	1.70	3.10
Koi	Xeon3.4	1.60	2.90
Koi	Xeon3.0	1.20	2.70
Koi	Xeon 3.06	1.10	2.30
Koi	Xeon 2.66	1.00	1.90

Fermilab Opteron Evaluation (Cont.)

- **Evaluation conclusions:**

  -x86_64 Linux is a stable operating platform

  -Opteron is a viable alternative to Xeon: yields slightly better performance with significantly less power draw at approximately the same price

Batch Systems Workshop

- **BQS@IN2P3**

  -BQS: Batch Queuing System

  -Custom software developed at IN2P3

  -Has support for managing AFS tokens

  -Scheduling policy based on resource requirements: BQS matches available resources on a worker with a job's requirements

  -Has support for parallel jobs

  -Runs on AIX, Solaris, and Linux

Batch Systems Workshop (Cont.)

- **FBSNG@Fermilab**
  - -FBSNG: Fermilab Batch System Next Generation
  - -Custom software developed at Fermilab
    - -No licensing issues
    - -Simple/lightweight design
  - -All Fermilab farms used FBSNG at one point
  - -Not suitable for use in a grid-environment
  - -Migrating to Condor

Batch Systems Workshop (Cont.)

- **IBM LoadLeveler**
  - -Version 3.3 available for Linux in 8/05
  - -Supports advanced node reservation: reservations can't overlap
    - -Administrator can impose reservation restrictions
- **LSF**
  - -Version 6.1 increases slot utilization from 74% (6.0) to 94%
  - -LSF multicluster in use at least 5 sites

Batch Systems Workshop (Cont.)

-Experimenting with grid metascheduling concepts and issues: Community Scheduler Framework (CSF)

-Several sites are using LSF itself as a metascheduler via a professional service contract with Platform: Platform is willing to help write LSF interfaces to other batch systems

-Working on ”Data Centric Scheduling”: dispatch jobs to machines where the cost of accessing requested data is least

Batch Systems Workshop (Cont.)

- Condor
  - -9 million lines of code
  - -Condor managing batch submissions on 10,000 systems at Micron
  - -Stork
    - -A scheduler for data placement
    - -Fault tolerant
    - -Can make space allocations and deallocations
    - -Supports many data transfer protocols (FTP, SRM, etc.)