Early Experience with Red Storm

April 19, 2005

Sue Kelly and Bob Ballance

Sandia National Laboratories

smkelly@sandia.gov, 505-845-9770

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Outline of Talk

- **Overview of Red Storm**
- **Where we are**
- **How we got here**
- **Where we're going**
- **What we learned**

NNSA National Nuclear Security Administration

Sandia National Laboratories

Red Storm is a Massively Parallel Processor

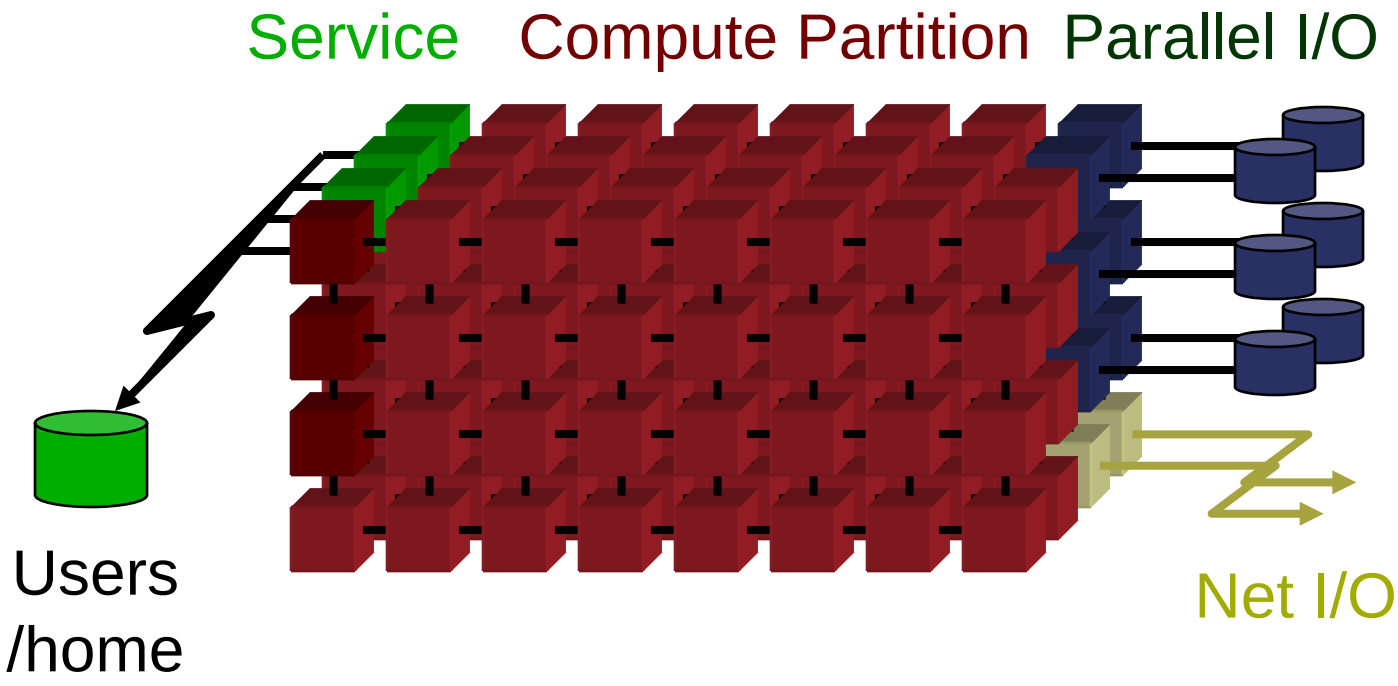

NNSA National Nuclear Security Administration

Sandia National Laboratories

27×16×24 3D Mesh/Torus + I/O

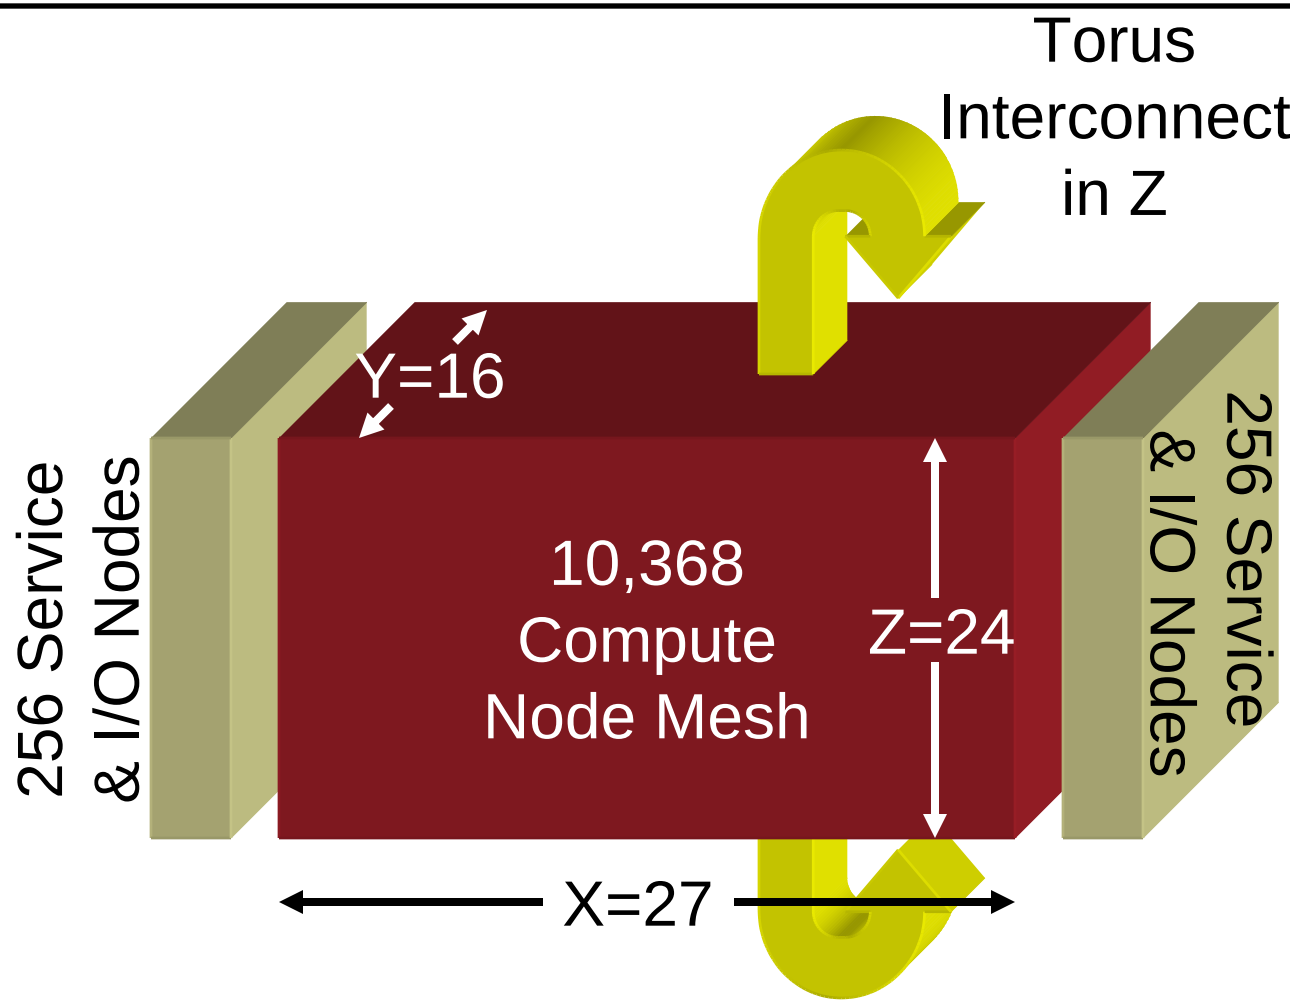

Usage Model

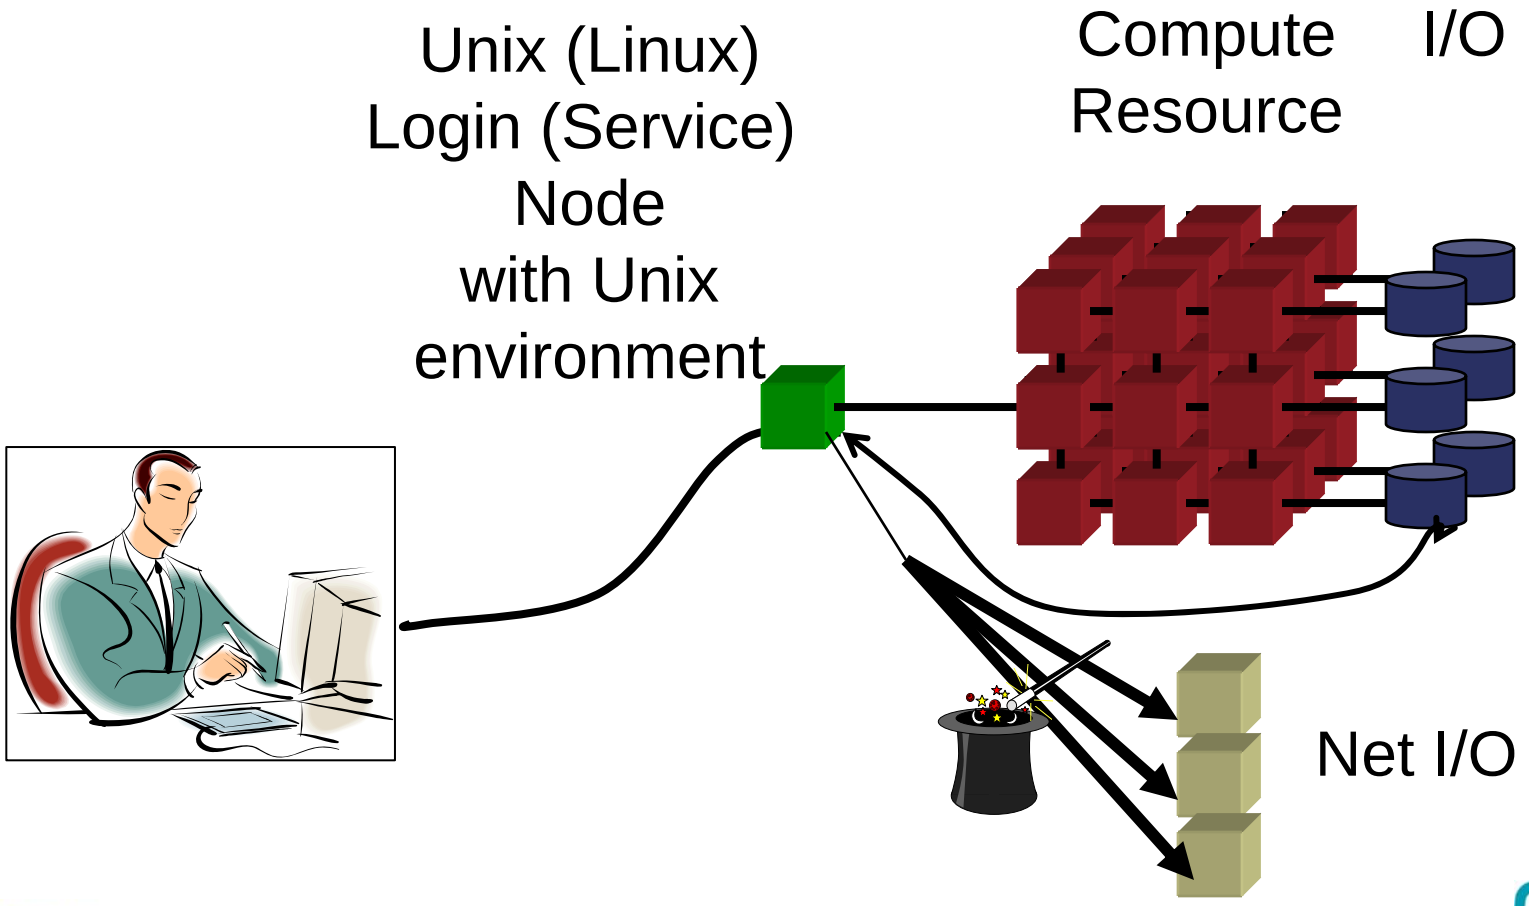
Unix (Linux)
Login (Service)
Node
with Unix
environment
Compute
Resource
I/O
Net I/O

Space Sharing of Jobs

- **Jobs occupy disjoint regions simultaneously**
- **Example – red, green, and blue jobs:**

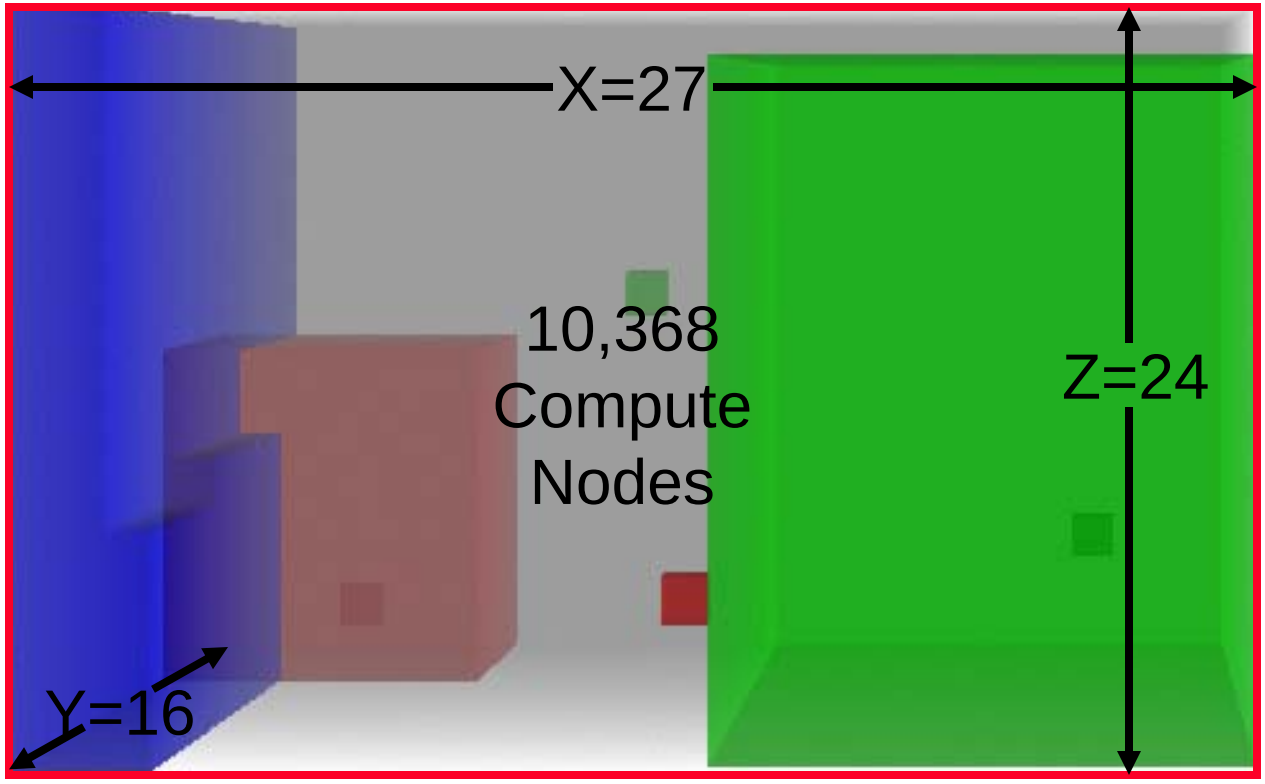

Red Storm Layout

(27 X 16 X 24 mesh)

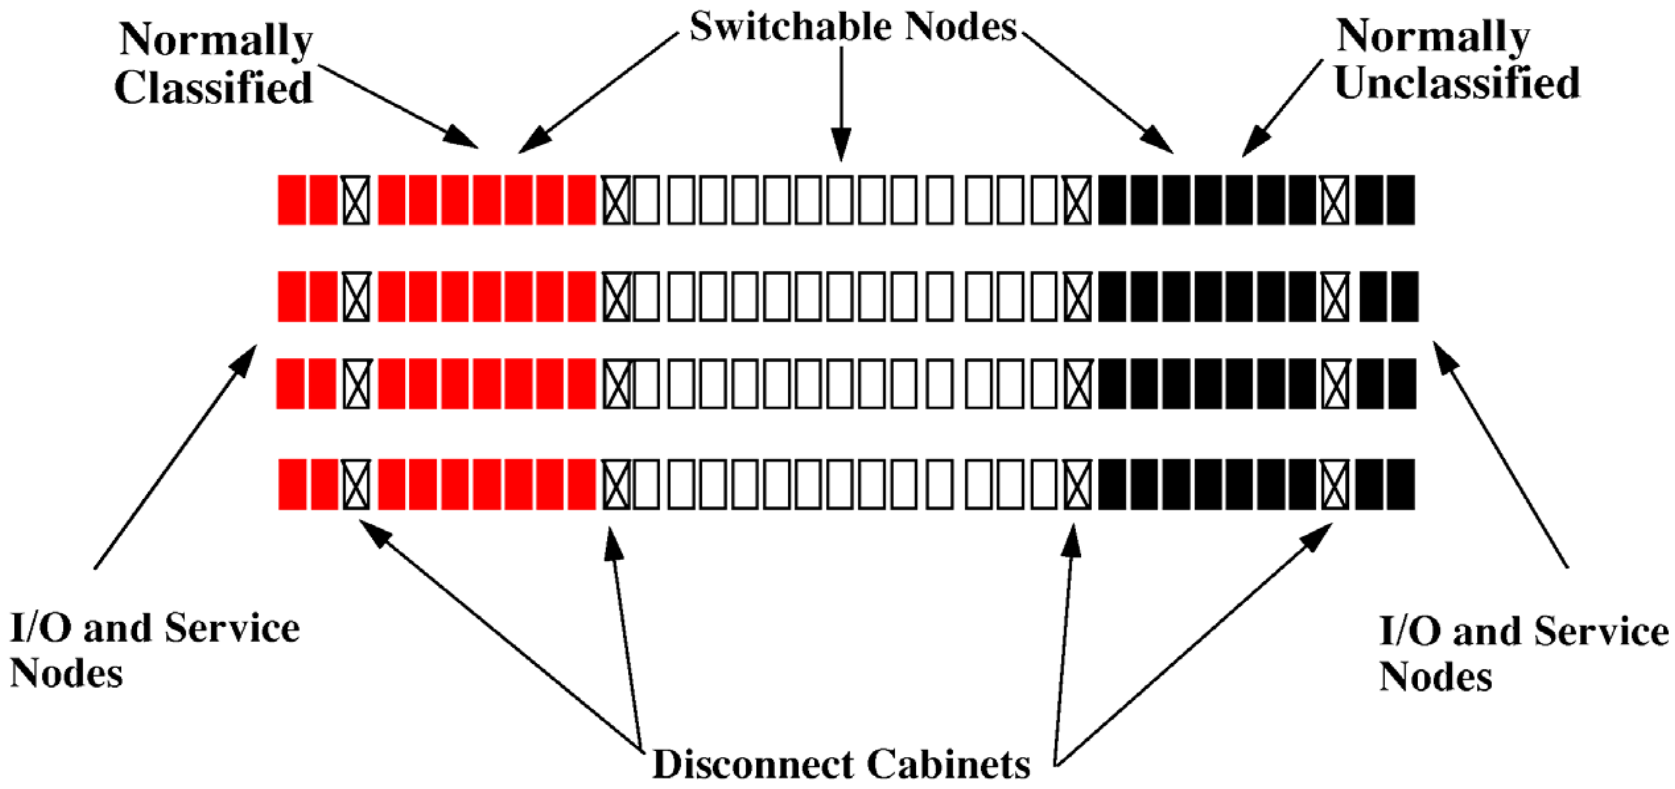

Disk storage system not shown

Red Storm Design Goals

Balanced System Performance - CPU, Memory, Interconnect, and I/O.

Usability - Functionality of hardware and software meets needs of users for Massively Parallel Computing.

Scalability - System Hardware and Software scale, single cabinet system to ~30,000 processor system.

Reliability - Machine stays up long enough between interrupts to make real progress on completing application run (at least 50 hours MTBI), requires full system RAS capability.

Upgradability - System can be upgraded with a processor swap and additional cabinets to 100T or greater.

Red/Black Switching - Capability to switch major portions of the machine between classified and unclassified computing environments.

Space, Power, Cooling - High density, low power system.

Price/Performance - Excellent performance per dollar, use high volume commodity parts where feasible.

Red Storm Design Parameters

True MPP, designed to be a single system.

Fully connected high performance 3-D mesh interconnect.

Topology - 27 X 16 X 24 compute nodes and 2 X 8 X 16 service and I/O nodes

108 compute node cabinets and 10,368 compute node processors. (AMD Sledgehammer @ 2.0 GHz)

~10 TB of DDR memory @ 333 MHz (1.0 GB per processor)

Red/Black switching - ~1/4, ~1/2, ~1/4.

8 Service and I/O cabinets on each end (256 processors for each color)

240 TB of disk storage (120 TB per color).

Red Storm Design Parameters

Functional hardware partitioning - service and I/O nodes, compute nodes, and RAS nodes.

Functional system software partitioning - LINUX on service and I/O nodes, LWK (Catamount) on compute nodes, stripped down LINUX on RAS nodes.

Separate RAS and system management network (Ethernet).

Router table based routing in the interconnect.

Less than 2 MW total power and cooling.

Less than 3,000 square feet of floor space.

Red Storm RAS System

RAS Workstations

- Separate and redundant RAS workstations for Red and Black ends of machine.
- System administration and monitoring interface.
- Error logging and monitoring for major system components including processors, memory, NIC/Router, power supplies, fans, disk controllers, and disks.

RAS Network - Dedicated Ethernet network for connecting RAS nodes to RAS workstations.

RAS Nodes

- One for each compute board
- One for each cabinet

Red Storm Performance

Peak of ~ 40 TF

Expected MP-Linpack performance >20 TF

Aggregate system memory bandwidth - ~55 TB/s

Interconnect

- **Aggregate sustained interconnect bandwidth > 100 TB/s**
- **MPI Latency - 2 μs neighbor, 5 μs across machine**
- **Bi-Section bandwidth ~2.3 TB/s**
- **Link bandwidth ~3.0 GB/s in each direction**

Disk and External Network I/O

- **Sustained 50 GB/s each color parallel disk I/O**
- **Sustained 25 GB/s each color external network I/O**

Comparison of ASCI Red and Red Storm

	ASCI Red	Red Storm
Full System Operational Time Frame	June 1997 (Processor and Memory Upgrade in 1999)	August 2004
Theoretical Peak (TF)	3.15	41.47
MP-Linpack Performance (TF)	2.379	>20 (est)
Architecture	Distributed Memory MIMD	Distributed Memory MIMD
Number of Compute Node Processors	9,460	10,368
Processor	Intel P II @ 333 MHz	AMD Opteron @ 2.0 GHz
Total Memory	1.2 TB	10.4 TB (up to 80 TB)
System Memory B/W	2.5 TB/s	55 TB/s
Disk Storage	12.5 TB	240 TB
Parallel File System B/W	1.0 GB/s each color	50.0 GB/s each color
External Network B/W	0.2 GB/s each color	25 GB/s each color
Interconnect Topology	3-D Mesh (x, y, z) 38 X 32 X 2	3-D Mesh (x, y, z) 27 X 16 X 24

Comparison of ASCI Red and Red Storm

	ASCI Red	Red Storm
Interconnect Performance MPI Latency Bi-Directional Link B/W Minimum Bi-section B/W	 15 µs 1 hop, 20 µs max 800 MB/s 51.2 GB/s	 2.0 µs 1 hop, 5 µs max 6.0 GB/s 2.3 TB/s
Full System RAS RAS Network RAS Processors	 10 Mbit Ethernet 1 for each 32 CPUs	 100 Mbit Ethernet 1 for each 4 CPUs
Operating System Compute Nodes Service and I/O Nodes RAS Nodes	 Cougar TOS (OSF1) VX-Works	 Catamount (Cougar) LINUX LINUX
Red Black Switching	2260 - 4940 - 2260	2688 - 4992 - 2688
System Foot Print	~2500 sq ft	~ 3000 sq ft
Power Requirement	850 KW	1.7 MW

Red Storm System Software

Operating Systems

- Linux on service and I/O nodes
- Catamount LWK on compute nodes
- Linux on RAS monitors

Run-Time System

- Logarithmic Job launch (yod)
- Node allocator (CPA)
- Batch system – PBS Pro

File Systems

- Lustre

User Environment

- PGI Compilers - Fortran, C, C++
- Libraries - MPI, I/O, Math, MPI-2
- Showmesh
- Debugger - TotalView
- Performance Monitoring

Network

- 50 x 10GigE to each RoSE cluster
- 10 x 1GigE to login nodes
- 1 GigE to Mgmt Workstations

System Mgmt and Admin

- Accounting
- Red Storm Management System

NNSA National Nuclear Security Administration

Sandia National Laboratories

Outline of Talk

- Overview of Red Storm
- Where we are
- How we got here
- Where we're going
- What we learned

NNSA National Nuclear Security Administration

Sandia National Laboratories

3/31 Milestone Achievements

1. All Red Storm hardware is on site, powered up, and integrated into the system.
2. We have booted 60 (3x20) compute cabinets as a single system and the remainder as two smaller systems. (Now booting as two systems total: 3x27 and 1x27.)
3. We have run all 7x applications.
4. We ran multiple applications simultaneously and demonstrated space sharing of the machine.
5. We have demonstrated sufficient stability of the platform and the software to verify normal operation of each application.
6. We used HPL (High Performance Linpack) to obtain a first look at the performance of the system.

Milestone #30: Bring Red Storm to initial operation and run the 7x applications

NNSA National Nuclear Security Administration

Sandia National Laboratories

We have run all 7x applications using a 7x dataset and normal I/O

Application	Problem Set	Processor Count	Run Time (min)
CTH	shape charge	2500 (5120)	60 (11)
Partisn	sntiming	2575	51
Partisn	sntiming 5x5	4096	Not reported
sPPM	inputdeck	2550	81
sPPM	inputdeck	3600	7
ITS	starsat	2370 (3500)	51 (2)
Sage	timing_c	2575	6
UMT2k	3d_slit	2575 (2799)	19 (1)
Salinas	Cube	2197	54
Alegra	ic_scaling	2048	205
Presto	walls	1024 (1536)	60 (30)
Calore	teratest	512 (1536)	119 (130)

Numbers in parentheses are runs with higher processor count, but with limited I/O or non 7x dataset

NNSA National Nuclear Security Administration

Sandia National Laboratories

Stability Metrics

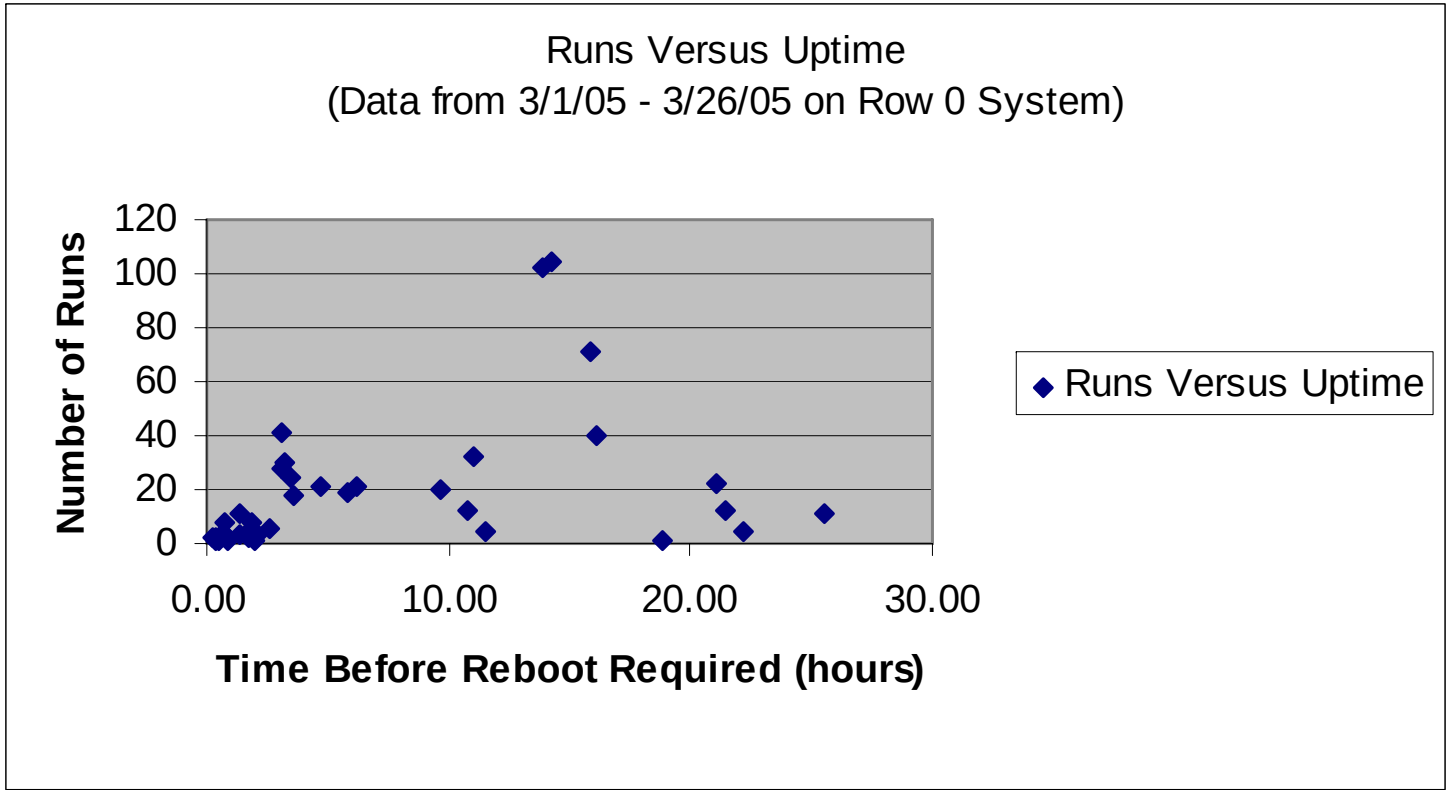

Points near origin are early attempts at larger configurations

NNSA National Nuclear Security Administration

Sandia National Laboratories

Outline of Talk

- Overview of Red Storm
- Where we are
- How we got here
- Where we're going
- What we learned

Condensed History of Red Storm Development

- Cray/Sandia partnership formed in July, 2002 with a letter contract
- Hardware arrived between September 2004 and February 2005
- Even though hardware had arrived, software was immature and had little exposure to real hardware
  - In particular RAS and networking software were deficient
- System was divided up into ~15 smaller systems for various development and risk mitigation efforts
- Many, many long days, nights, and weekends got us to 3/31

NNSA National Nuclear Security Administration

Sandia National Laboratories

Sample Application Scaling Progress From Row 0 Runs

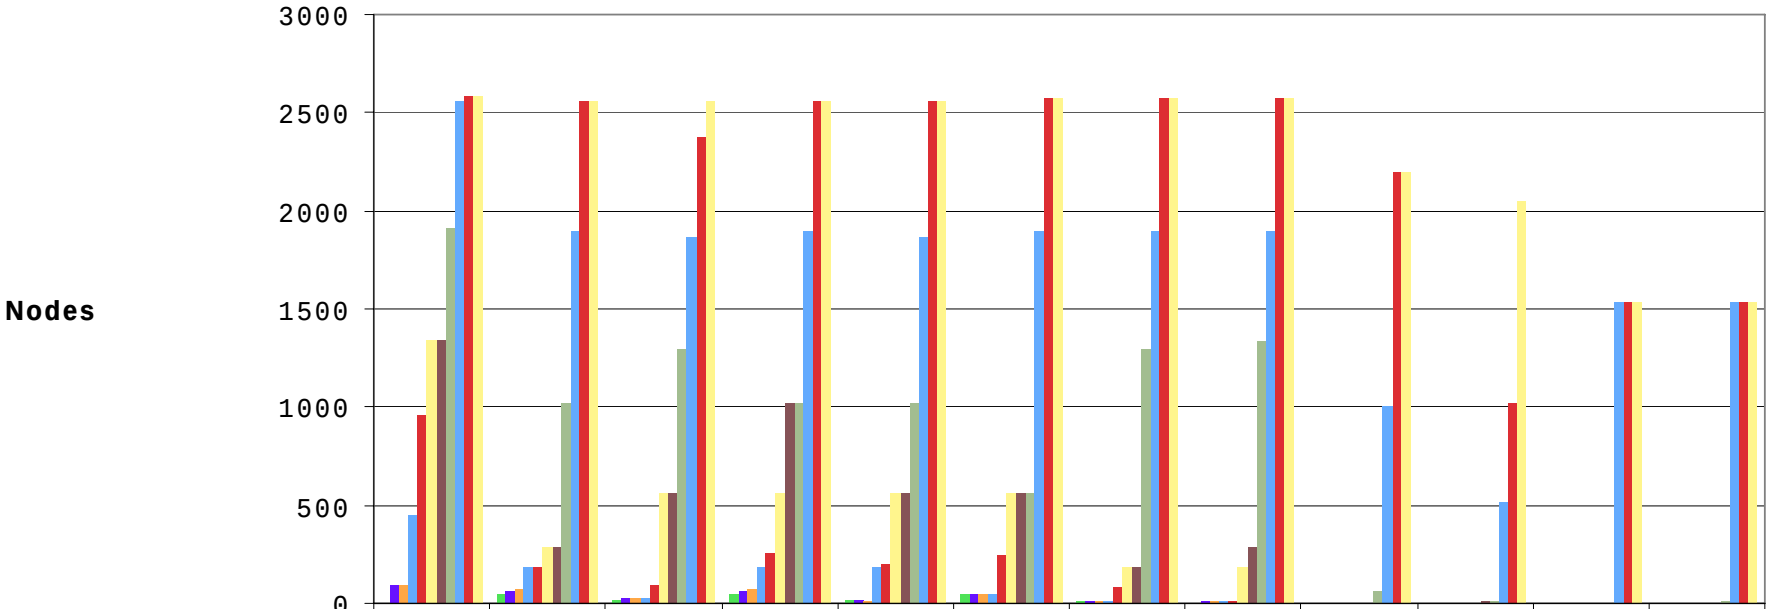

	Boot	HPL	sPPM	CTH	ITS	Partisn	Sage	UMT2k	Salinas	Alegra	Calore	Presto
1/21/2005 Sandia		44	16	44	16	44	5					
1/28/2005 Sandia	96	64	24	64	16	44	5	5				
2/04/2005 Sandia	96	70	24	70	7	48	5	5				
2/11/2005 Sandia	448	188	24	188	188	48	5	6				
2/18/2005 Sandia	960	188	92	256	200	252	80	6				
2/25/2005 Sandia	1344	284	560	564	564	564	188	188				
3/04/2005 Sandia	1344	284	560	1024	564	564	188	284	0	8		
3/11/2005 Sandia	1915	1024	1300	1024	1024	564	1300	1331	64	8	4	8
3/18/2005 Sandia	2562	1900	1872	1900	1872	1900	1900	1900	1000	512	1536	1536
3/24/2005 Sandia	2582	2562	2375	2562	2560	2575	2575	2575	2197	1024	1536	1536
3/31/2005 Sandia	2582	2562	2560	2562	2560	2575	2575	2575	2197	2048	1536	1536

Applications

NNSA National Nuclear Security Administration

Sandia National Laboratories

Latency performance over time

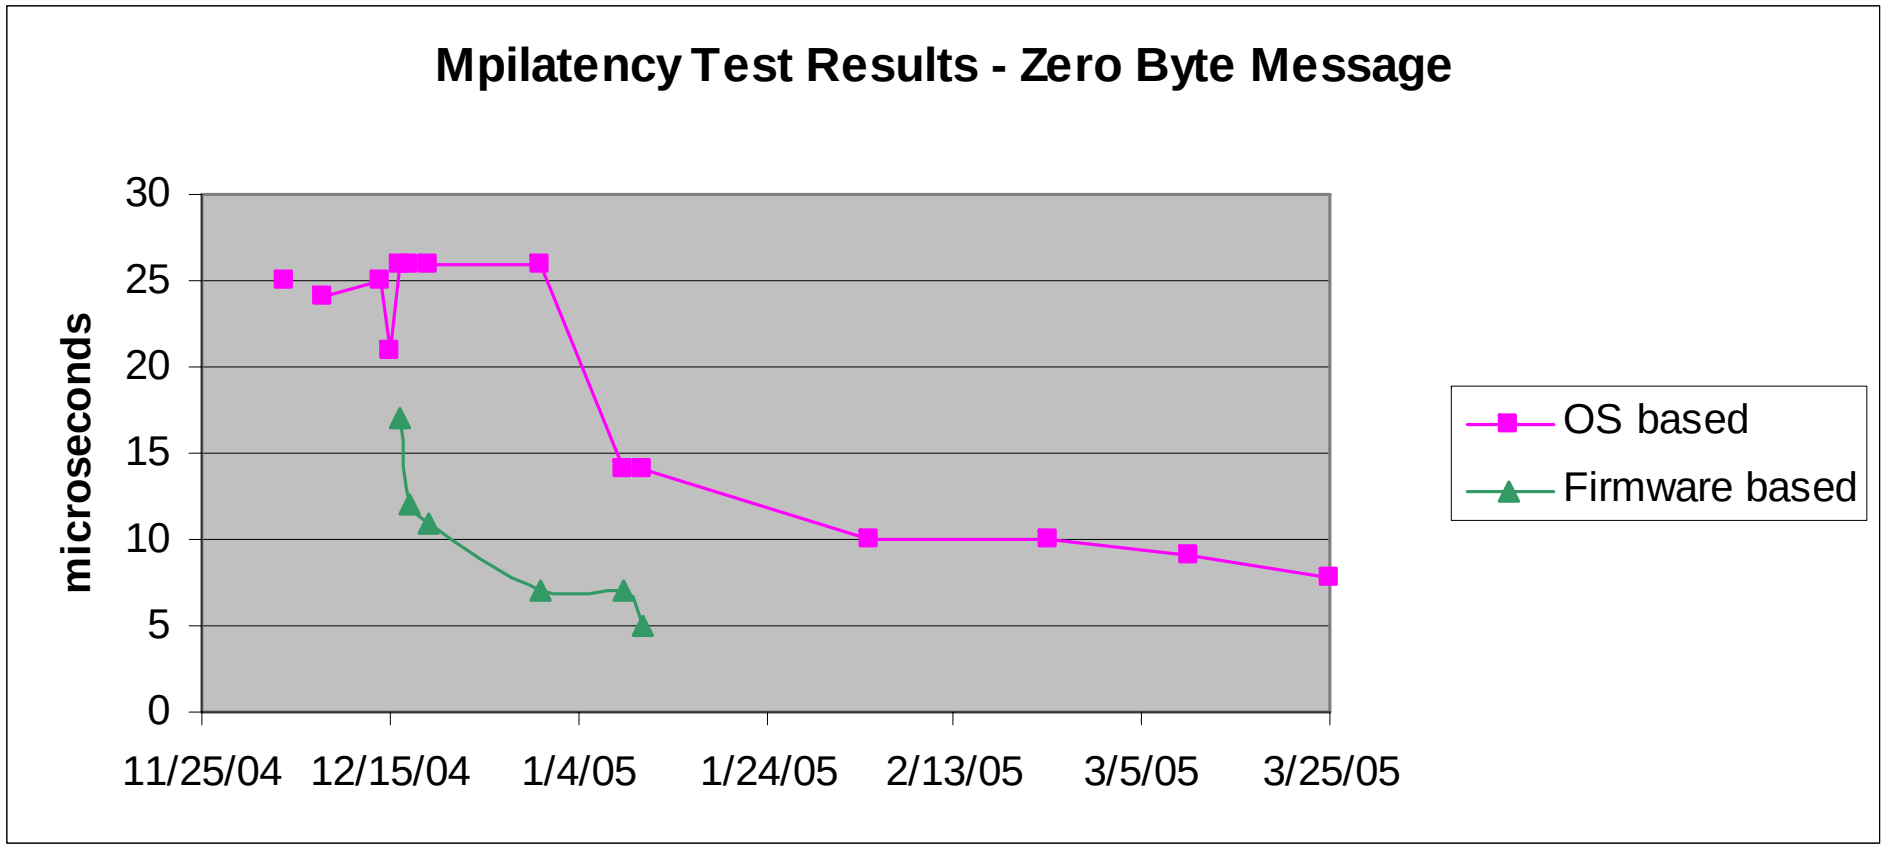

NNSA National Nuclear Security Administration

Sandia National Laboratories

Outline of Talk

- Overview of Red Storm
- Where we are
- How we got here
- Where we're going
- What we learned

NNSA National Nuclear Security Administration

Sandia National Laboratories

Upcoming Milestones

- **Begin classified operation in June for LANL runs only**
- **Complete 3 ASC milestones by September 05, which will result in a production-ready system for a limited set of users**

Short Term Plans

- Service partitions need to move beyond:
  - Minimal number of nodes
  - Minimal configurations
  - Not operating as intended
  - PBS/Job batch policies not installed
  - No normal login nodes and login procedures
  - Very little Sandia network integration
- Lustre needs more testing
- Need more reliability testing

Outline of Talk

- Overview of Red Storm
- Where we are
- How we got here
- Where we're going
- What we learned

Lessons Learned

- Seek first to emulate
  - Learn from the past
  - Simulate the future
- Successful technology transitions require people transformations
  - Need technology philosophers
    - Tilt Meters
    - Historians
- Even Tiger Woods has a coach
- The big bang only worked once
  - Deploy test platforms early and often
- Build de-scalable, scalable systems
  - Don't forget that you have to get it running first!
  - Leave the support structures (even non-scalable development tools) in working condition, you'll need to debug some day
- Only dead systems never change
  - Nobody ever built just one system even when successfully deploying just one system
  - Nothing is ever done just once
- End to end arguments apply within large systems AND teams
- Build scaffolding that meets the structure
  - Is build and test infrastructure in place FIRST?
  - Will it effectively support both the team and the project?

Conclusion

- Red Storm is on target to be the 7X Son of Red
- Early experience shows good scalability
- We have much to do in the next 5 months

NNSA National Nuclear Security Administration

Sandia National Laboratories