SAND2008-3565P

The Sandia HPC Story

Robert A. Ballance

28 May 2008

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.

Sandia HPC History

- **1987 – 1024 node nCUBE 10**
  - Sandia fielded first true MPP
  - Won Karp Challenge
  - Won inaugural Gordon-Bell Prize
  - Awarded patents for nearly every aspect of MPP software
- **1990 – 2 1024-node nCUBE-2's**
  - Outperformed Cray vector computers @ 1/7 the cost!
  - Sandia began research on Light Weight Kernel (LWK) operating systems (SUNMOS)

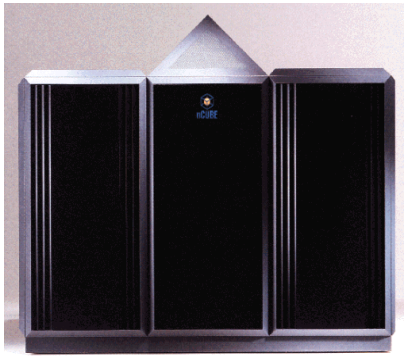

- **1993– Sandia fielded first ~1850-node (3900 processors) Intel Paragon**
  - –Sandia (Intel) Paragon is #1 on Top 500 list (first for Sandia)
  - –Wins Sandia's second Gordon-Bell Prize
  - –First use of Sandia developed Light Weight Kernel (LWK) Operating System software (PUMA) for production computing (16 MB memory/node.)

Sandia National Laboratories

1997– Tflops, world's first Terascale computer

4600+ nodes (9200+ processors – initial 200 MHz, upgraded to 333 MHz; memory doubled)

- –Ran Sandia system developed LWK System Software (Cougar)
- –Number 1 on the "Top 500" list for 7 consecutive lists from June, 1997 through June, 2000, a record still unmatched! **(1.068, 1.338, 2.123, 2.3796 TF)**
- –Red-Black switching
- –Approx 16 Cray YMP's
- –Portals networking protocol

- ASCI Option Red – The Intel Tflops system was ***operational until October of 2005***. Subsequently retired and reduced to rubble for security reasons.

The original Berkeley Sun NOW – 10 GF Nov 97 TOP500

Sandia HPC history

- 1997– Sandia began development of world's first Linux "Super-cluster", Cplant
  - Sandia integrated DEC/Compaq HW with Myrinet high performance interconnect network
  - Sandia developed all run-time, file system and messaging software
  - Open source Cluster Management Toolkit
  - World's first terascale cluster– eventually achieves nearly 1 TF on Linpack. **(996.9 GF in November 2003)**
- 2001– Sandia invented Red Storm architecture – issues RFQ to industry
  - Responses were being reviewed on Sep 11
- 2002 – Sandia contracted with Cray, Inc to jointly develop Red Storm/XT3
  - Sandia architecture; Cray built hardware; Sandia and Cray developed system software

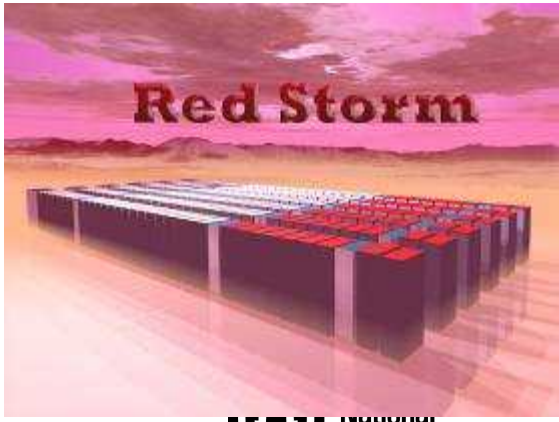

Red Storm

True MPP, designed to be a single system
- Full 3-D mesh interconnect
- 12,960 compute nodes
  - AMD Dual-Core Opterons @ 2.4GHz
  - 25,920 cores
- 39 Terabytes of memory
- 340 Terabytes of disk storage
- 3715 embedded RAS processors

Sandia contributions include
- MPP system architecture
- Helped design interconnect
- Lightweight kernel strategy
- Red/Black switching

Excellent performance
- 102.2 TeraOps on HPL
- Second system ever to exceed 100TF
- First general-purpose system to exceed 100TF

Successful technology transition
- Cray now has now > over 30 systems based the Red Storm model

Serves LANL, LLNL, SNL, Alliance Partners
- Remote user support a given
- Remote networking support active since ASCI Red

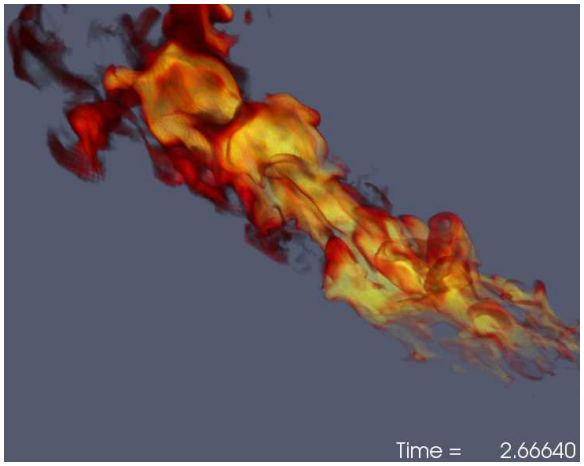

Fire simulations are being used to certify new test facility and will be used to analyze weapons safety issues

Red Storm: Major Events

- 2004 – 2005 Delivery & Integration
  - Nominal peak: 41 TF
- 2006 – Upgrade to dual core, 5th Row
  - Nominal peak: 124 TF
- 2008 –
  - 1.5PB disk,
  - 48% Quad Core, 2G/core (≈ 75 TB total)
  - Nominal peak: 284 TF
  - Starts next month!

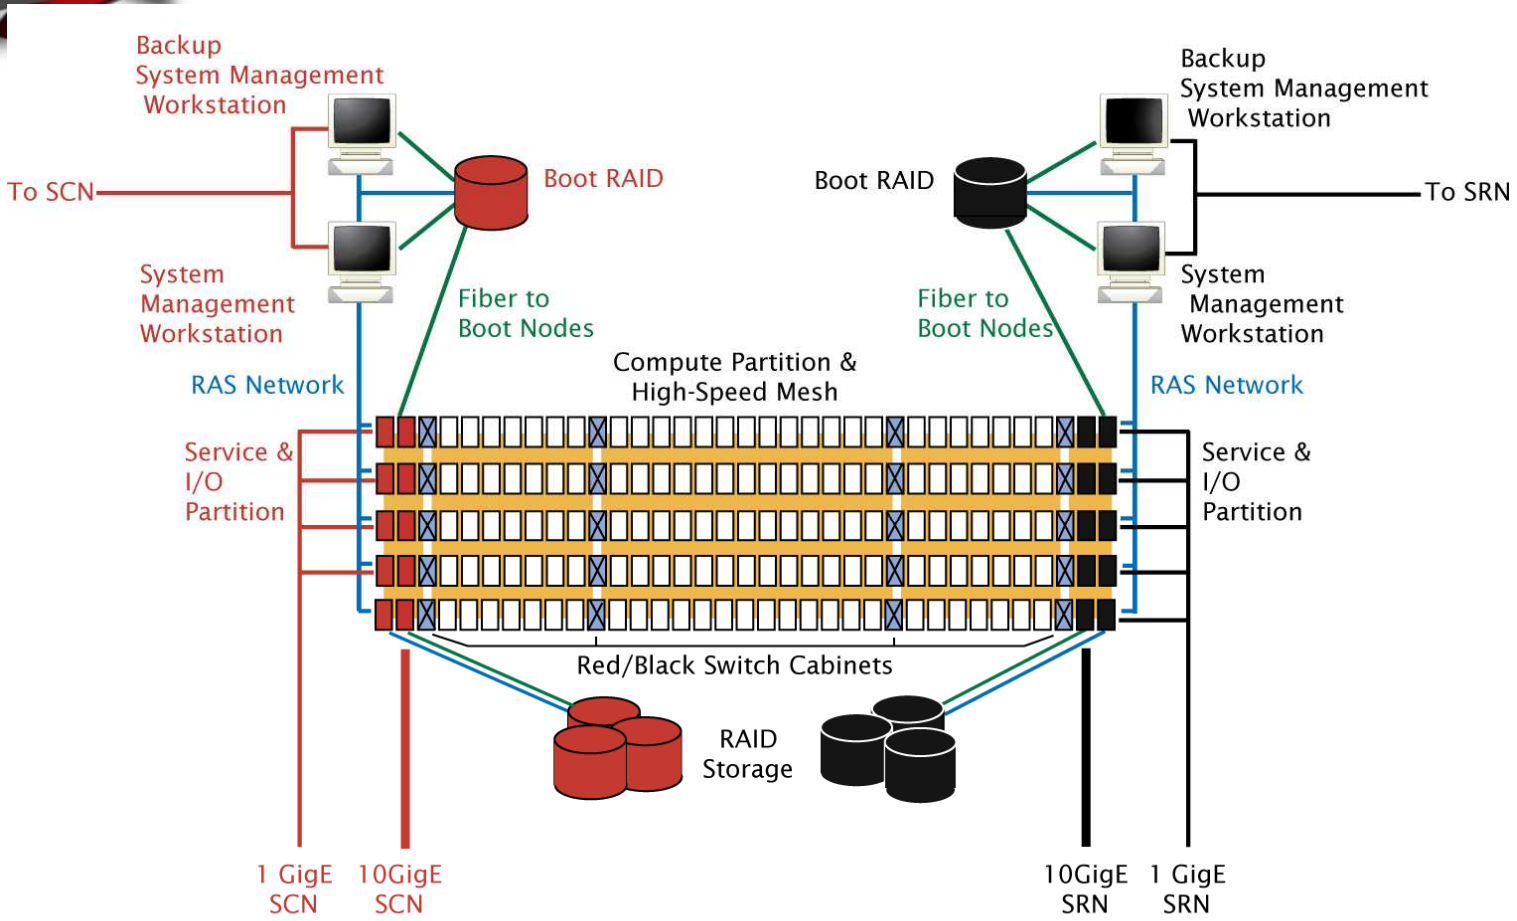

Red-Black Switching

- Red Storm was designed from the start to support red/black switching
- Switchable components have no persistent state that can be written from the Classified side

• Switch Cabinets isolate sections of the machine, so that the service sections remain in production at all times

• Network infrastructure is balanced between black and red

- Design based on ASCI Red
  - Familiar to NNSA
  - Changeover O(1hr)
- RAS network is separable
- ASCI Red could switch without reboot
  - So far, Red Storm is unable to achieve this
- Advantages
  - Load balancing of work
  - Flexibility during upgrades and mtce
  - Low incremental cost

Red Storm – System Software

- Operating Systems
  - Linux on service and I/O nodes (SuSE Enterprise Server 9.0)
  - Catamount lightweight kernel on compute nodes
  - Linux on RAS monitors
- Run-Time System
  - Logarithmic job launch (yod)
  - Node allocator (CPA)
  - Batch system – PBS Pro/MOAB
- File Systems
  - High performance file system (lustre)
- User Environment
  - PGI compilers – Fortran, C, C++
  - Libraries – MPI, I/O, Math, MPI-2
  - Showmesh
  - Debugger – Totalview
  - Performance Monitor
- Network
  - 50 x 10 GigE to RoSE, WAN
  - 10 x 1 GigE to login nodes
  - 1 GigE to Mgmt stations
- System Mgmt and Admin
  - Accounting
  - Red Storm Management System

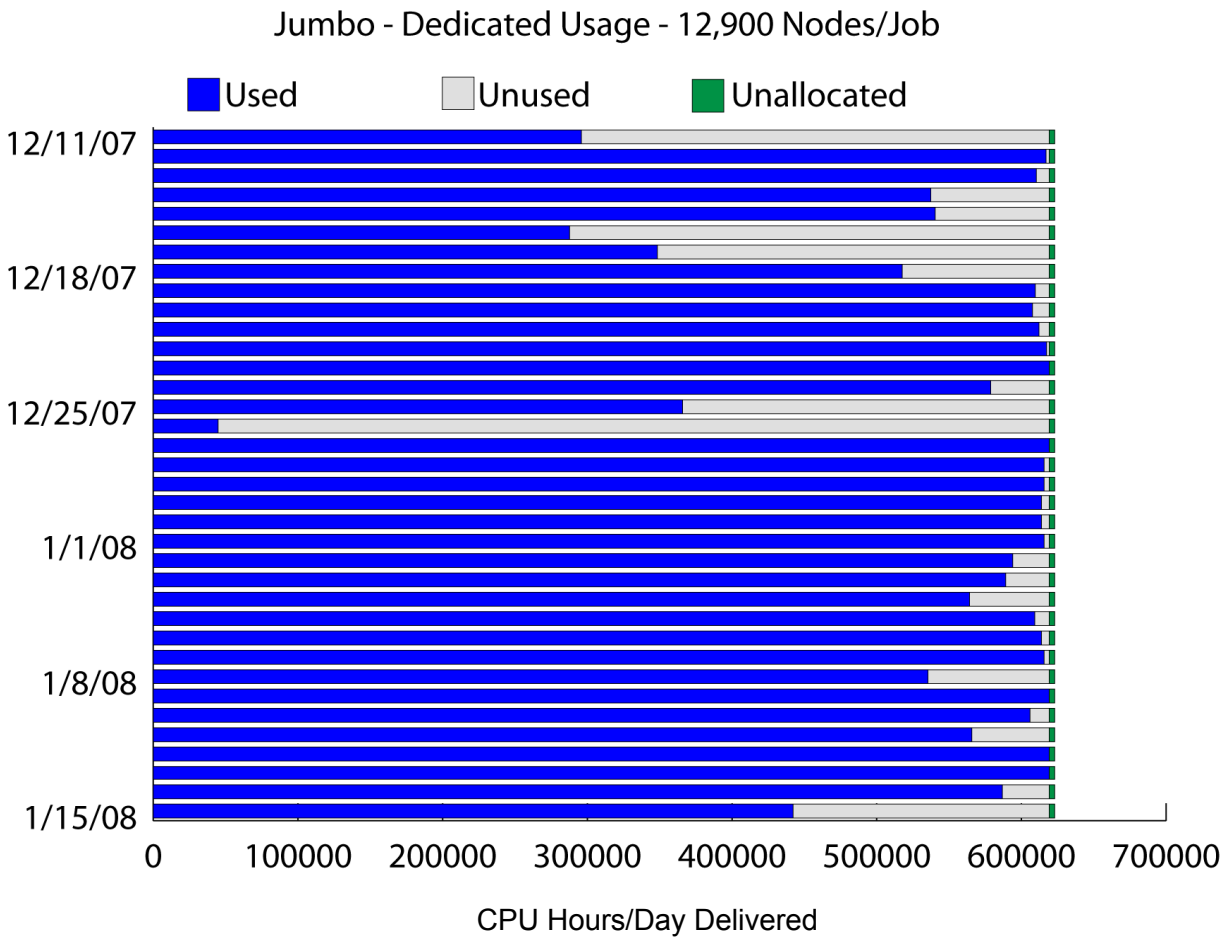
Jumbo - Dedicated Usage - 12,900 Nodes/Job
Used
Unused
Unallocated
12/11/07
12/18/07
12/25/07
1/1/08
1/8/08
1/15/08
0
100000
200000
300000
400000
500000
600000
700000
CPU Hours/Day Delivered

Architected Red Storm Environment

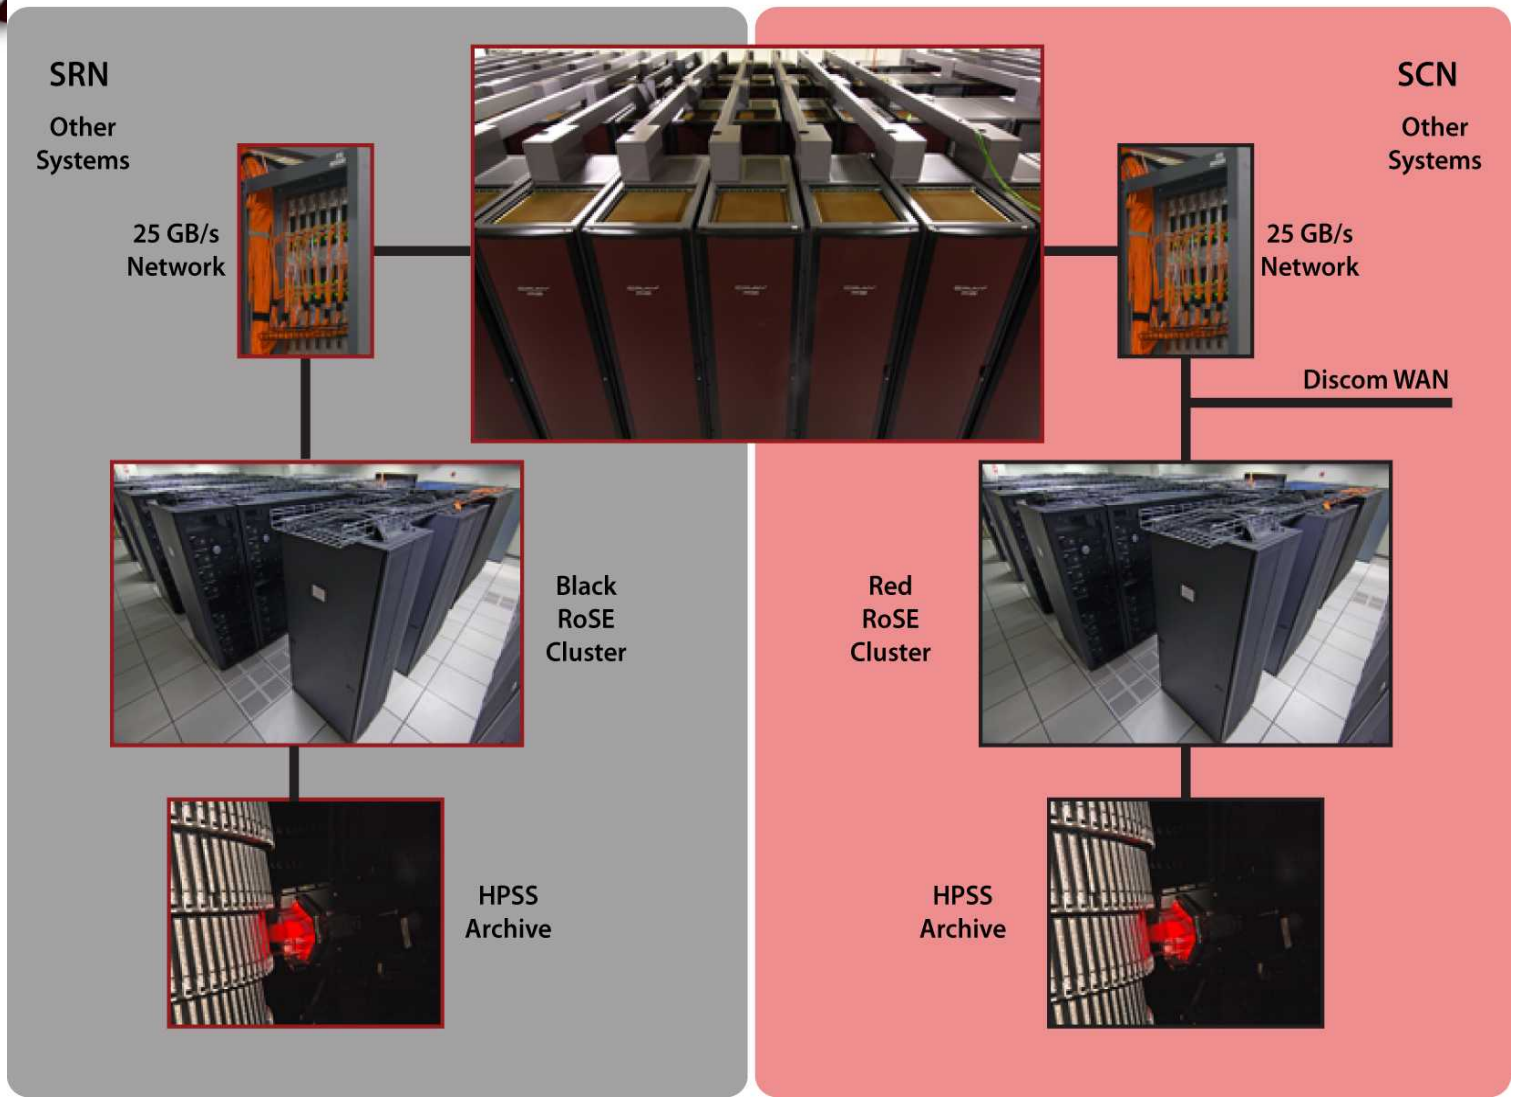

Current CapViz/ESHPC Resources (5-2008)

System	Network	Nodes/Processor cores	Processor	Memory	TFLOPS	Processor Hours per Year	GA Date
Tbird	SRN	4480 nodes/8960 cores	3.6 GHz EM64T	6 GB/node	65	78,489,600	Mar-2006
Spirit	SRN (NW)	512 nodes/1024 cores	3.4 GHz EM64T	2 GB/node	6.9	8,970,240	Nov-2004
Shasta	SRN (CA)	198 nodes/396 cores	3.0 GHz IA32	2 GB/node	2.4	3,468,960	Nov-2003
Black RoSE	SRN (Viz)	128 nodes/512 cores	2.8 GHz AMD	8 GB/node	2.8	4,485,120	Mar-2007
Rogue	SRN (ESHPC)	407 nodes/814 cores	2.8 GHz EM64T	4 GB/node	4.6	7,130,640	
Buzzard	SRN (ESHPC)	150 workstations	2.2 to 3.02 GH Xeon and EM64T	1 to 4 GB		131,400	
Razor	SCN (ESHPC)	280 nodes/1120 cores	3.0 GHz Woodcrest	8 GB/node	13.4	9,811,200	Jan-2007
Justice	SCN (NW)	512 nodes/1024 cores	3.4 GHz EM64T	2 GB/node	6.9	8,970,240	Nov-2004
Lassen	SCN (CA NW)	192 nodes/384 cores	3.4 GHz EM64T	2 GB/node	2.6	3,363,840	Dec-2004
Red RoSE	SCN (Viz)	264 nodes/528 cores	3.6 GHz Xeon	4 GB/node	3.8	4,625,280	June-2006
Unity	*SCN (NW)*	*272 nodes/4352cores*	*2.2 GHz AMD Quad Core*	*32 GB/node*	*38*	*38.123.520*	*Jul-2008*
Glory	*SCN (NW)*	*272 nodes/4352 cores*	*2.2 GHz AMD Quad Core*	*32 GB/node*	*38*	*38.123.520*	*Aug-2008*
Whitney	*SCN (CA NW)*	*272 nodes/4352 cores*	*2.2 GHz AMD Quad Core*	*32 GB/node*	*38*	*38.123.520*	*Sep-2008*

SRN Multi-Cluster file systems:
350 TB Lustre (/gscratch1)
50 TB Panasas (/gscratch3)

SCN Multi-Cluster file systems:
1000 TB Lustre (/gscratch1)

Sandia National Laboratories

Thunderbird: 2005 - Present

- 4480 Dell PE1850 Servers
- Infiniband
- 8960 Intel 3.6 GHz EM64T processors
- 64.51 Tflops theoretical peak
- Online in 2005
- 2006: TBird becomes the largest cluster of its type in the world using Open Fabrics / Open MPI.

TLCC: 2008

- Tri-Labs Purchase of Capacity Systems
- 20 TF/scalable unit (288 nodes, 4 x 4 cores/node)
- Currently deploying 4 SU (80 TF)

SNL and the History of Lightweight Kernels

- 1993 SUNMOS (Sandia-UNM Operating System) on Paragon
- 1995 - Puma – Rewrite of SUNMOS as part of a research project
- 1997 - Cougar – Port of Puma running ASCI Red.
  - Portals protocol developed by SNL + Univ of New Mexico
- 2003 - IBM uses LWK concepts from Cougar for the Blue Gene kernel
- 2004 - Catamount – Port of Cougar to run on Red Storm
- 2006 - Catamount VN – for dual core processors
- 2007 - Research - N-way Catamount for multicore
- 2008 - Research - Virtualization Layers


725

Advanced Simulation & Computing™

Background

Sandia National Laboratories

Red Storm Architectural Goals

- Balanced System Performance: CPU, Memory, Interconnect and I/O
- Scalability: System Hardware and System Software scale, a single cabinet system to 32K processor system
- Functional Partitioning: Hardware and System Software
- Reliability: Full system Reliability, Availability, Serviceability (RAS) designed into Architecture
- Upgrade-ability: Designed in path for system upgrade
- Red/Black Switching: Flexible support for both classified and unclassified Computing in a single system
- Custom Packaging: High density, relatively low power system
- Price/Performance: Excellent performance per dollar, use high volume commodity parts where feasible

Red Storm – Lightweight Kernel

- Lightweight Kernels are a key contribution from Sandia to the HPC community
- Based on several key principles:
  - The purpose of compute nodes is to compute
    - Whenever any processor takes an interrupt, other processors in the computation may have to wait
    - Your computation is only as fast as its slowest component
  - Maximize the resources devoted to the computation
  - Simplify the code to enhance reliability
- This may not be an issue for dozens of nodes
  - Red Storm has 25,920 processors to coordinate!

Red Storm – Lightweight Kernel

- A lightweight compute node OS is fundamental to the Sandia architecture
- It is essential for:
  - Maximizing CPU resources
    - Reduce OS and runtime system overhead
  - Maximizing memory resources
    - Small memory footprint, large page support
  - Maximizing network resources
    - No virtual memory, physically contiguous address mapping
  - Increasing reliability
    - Small code base, reduced complexity
  - Deterministic performance
    - Repeatability
  - Scalability
    - OS resources must be independent of job size