



Operations at Scale; Lessons to be Remembered

Robert A. Ballance, raballa@sandia.gov
John P. Noe, jpnoe@sandia.gov

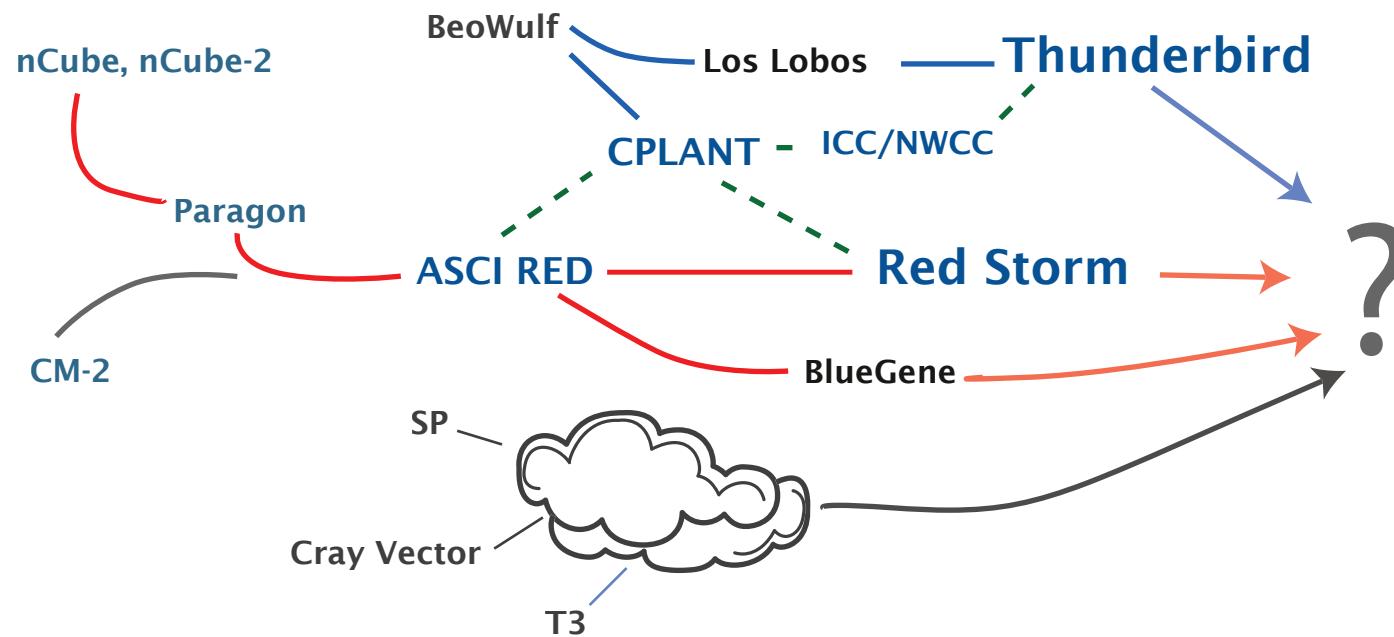
14 May 02007

SAND 2007-XXXX

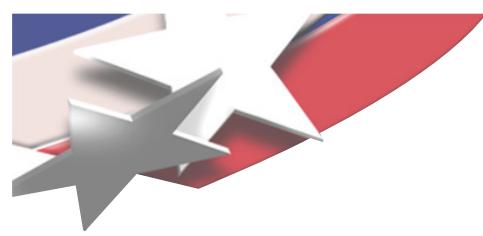


Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company
for the United States Department of Energy's National Nuclear Security Administration under
contract DE-AC04-94AL85000.

Map of Talk: Hardware Projection



Whirlwind • Stretch • BBN • ETA • Alliant • Elexsi • MasPar • Convex • Kendall Square • CDC • Cray 3 • ...



Red Storm & Thunderbird

	Red Storm	Thunderbird
Vendor	Cray/Sandia	Dell
Nov 2006 Rank	#2	#6
Stance	Classified + Unclassified	Unclassified
Service Nodes	320 + 320	
Processor	2.4 Ghz Dual-core Opteron	3.6 Ghz EMT64
Compute Nodes	12,960	4,480
Compute Cores	25,960	8,960
Segments	3360/6240/3360	5
Interconnect	SeaStar Mesh	Infiniband
Disk	170TB + 170TB	46TB + 420TB shared
OS	Linux + Catamount	Linux
Job Size (Cores)	> 1024, many > 5,000	< 1024 PE

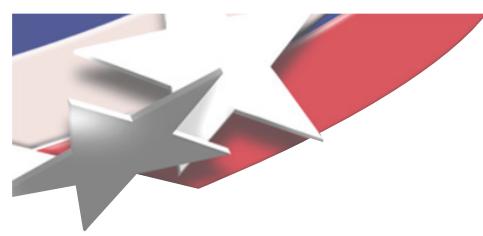


Map of Talk: Conceptual Projection



“I want to build a clock that ticks once a year. The century hand advances once every one hundred years, and the cuckoo comes out on the millennium. I want the cuckoo to come out every millennium for the next 10,000 years. If I hurry I should finish the clock in time to see the cuckoo come out for the first time.”

Danny Hillis





Seek First to Emulate

A complex system that works is invariably found to have evolved from a simple system that works. *John Gall*

- **Learn from the past**
 - The role of failure in system (bridge) design
 - Sibley's 30-year cycle
- **Simulate the future**
 - Systems are too large to start fixing after they are built
 - One of the first things a computer does to design the next computer!

The big bang only worked once

- **Nobody ever builds just one system**
 - Single systems change over time
 - Need for consistency checking
 - Prototypes!
- **Globalize agility; localize fragility**
- **Deploy test platforms early and often**
 - System test
 - Software checkout
 - Application test
 - Regression testing
- **Only dead systems never change**
 - Livable systems are automated
 - Living systems get smarter over time
 - Teams can get smarter over time



Hubble Image of NGC 2440



Build descachable scalable systems

- Scalability has to be designed into the system from the start
 - Even small details can hurt you; the Alegra story
- Never forget that you have to get it running first
 - Argument: We can't add logging; it will slow down the system.
 - Build scaffolding that meets the structure
- Is the build/test/benchmark infrastructure in place first?
 - Will it effectively support the installation team, the users, and operations?
- Leave the support structures (even non-scalable ones) in working condition
 - You'll need to debug someday
 - Like yesterday!
 - This means you have to test the testers





When the lights turn green, better recheck the connections!

- **Software only reports reality as it sees it**
 - You can't really trust software when it is new
 - You might be able to trust it after considerable use
 - You can't ever trust software that believes itself
- **Requirements for management software**
 - *Explore* to see what is out there, and make that information part of the internal view.
 - *Coerce* what is out there to match the internal view
 - *Compare* internal structures and the external reality
 - Depth perception!
- **Parallel tools for parallel systems**

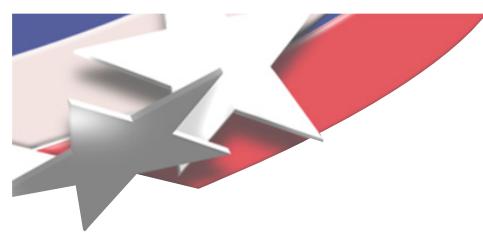
"Ground truth in Reality"

Jan Newton

End-to-End Arguments Apply

Building complex function into a network implicitly optimizes the network for one set of uses while substantially increasing the cost of a set of potentially valuable uses that may be unknown or unpredictable at design time

- **Within large systems**
 - complexity at edge
 - flexibility at core
- **Within teams**
 - communication structures
 - decision-making structures



Even Tiger Woods has a coach

- **Don't assume you know/understand it all**
 - Observers help
 - Open processes
 - Transparency: of Process, Code, and Operations
 - Collaborative systems
- **Never underestimate your blind spots**

Play with your mental blocks!

Risk Analysis (Kaplan & Garrick, 1981)

- *What can go wrong?*
- *How likely is it to happen?*
- *What are the consequences?*

Add a fourth:

- *How will we know it has happened?*



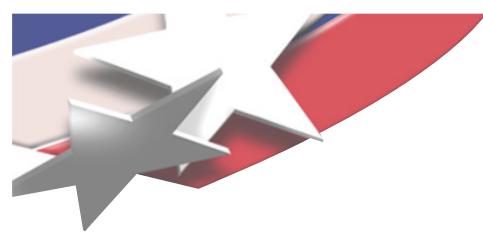
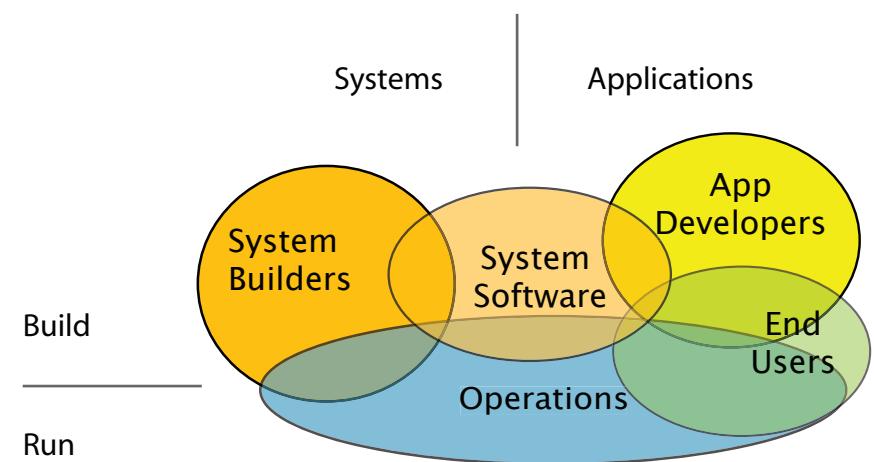


Successful technology transitions require people transformations

- Roles for veterans
 - Philosophers
 - Tilt meters
 - Historians
 - The Bell Labs experience
- What is the right ratio of veterans to newbies?
 - 1:5?
 - 1:10?
 - 5:1?

Begin with the End in Mind

- Involve Operations from Day 1
- Making it work cannot be a downstream task
- Operations folks are scouts
 - They'll figure out how to make it work
 - They probably understand the terrain and the natives



Mind the Long Term

- **Trust the future**

- There will be a next system
- Beware Brooks' 2nd system effect!



What is the Apgar score?

"One minute — and again five minutes — after your baby is born, doctors calculate his Apgar score to see how he's doing. It's a simple process that helps determine whether your newborn is ready to meet the world without additional medical assistance.

This score — developed by anesthesiologist Virginia Apgar in 1952 and now used in modern hospitals worldwide — rates a baby's appearance, pulse, responsiveness, muscle activity, and breathing with a number between zero and 2 (2 being the strongest rating). The numbers are totaled, and 10 is considered a perfect score." [2]

- **Measure for life**

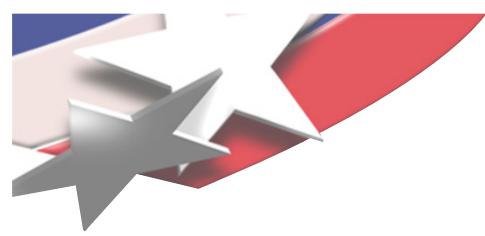
- LINPACK as Apgar
- What is the HPC equivalent of a lifetime achievement award?
- NERSC ESP Benchmark is one inspiration



THE HONORARY AWARD (Statuette). *This award shall be given to honor extraordinary distinction in lifetime achievement, exceptional contributions to the state of motion picture arts and sciences, or for outstanding service to the Academy.*

Restatement

- Seek first to emulate
- The big bang only worked once
- Build descalable scalable systems
- Make the lights green, then recheck the connections
- Even Tiger Woods has a coach
- End-to-end arguments apply
- Successful technology transitions require people transformations
- Begin with the end in mind
- Mind the long term





Design Principles: Clock of the Long Now

- **Longevity:** Display the correct time for ten millennia.
- **Maintainability:** with Bronze-age technology if need be.
- **Transparency:** obvious operational principles.
- **Evolvability:** improvable over time.
- **Scalability:** the same design should work from tabletop to monument size.

From The Clock of the Long Now, Stewart Brand, 01999



Acknowledgements

- **For patience, good humor, new opportunities, elegant insights, and quotable phrases:**
 - Ron Brightwell, Bill Camp, Sophia Corwell, Joe Davison, Frank Gilfeather, Michael Hannah, Jim Harrell, Tram Hudson, Steve Johnson, Sue Kelly, Ruth Klundt, Jim Laros, Rob Leland, Barney Maccabe, Michael Mahon, Geoff McGirt, John Naegle, Kevin Pedretti, Rolf Riesen, Brian Smith, Jon Stearley, Jim Sundet, Jim Tomkins, Michael Van De Vanter, John Van Dyke, David Wallace, and Lee Ward
- **Finally, thanks to all the system administrators who are keeping our systems running, even as we speak.**

