



Testing & Integration of HEC Systems at NCAR

Petascale Systems Integration into Large Scale Facilities Workshop
Hotel Nikko, San Francisco, CA
15 May 2007

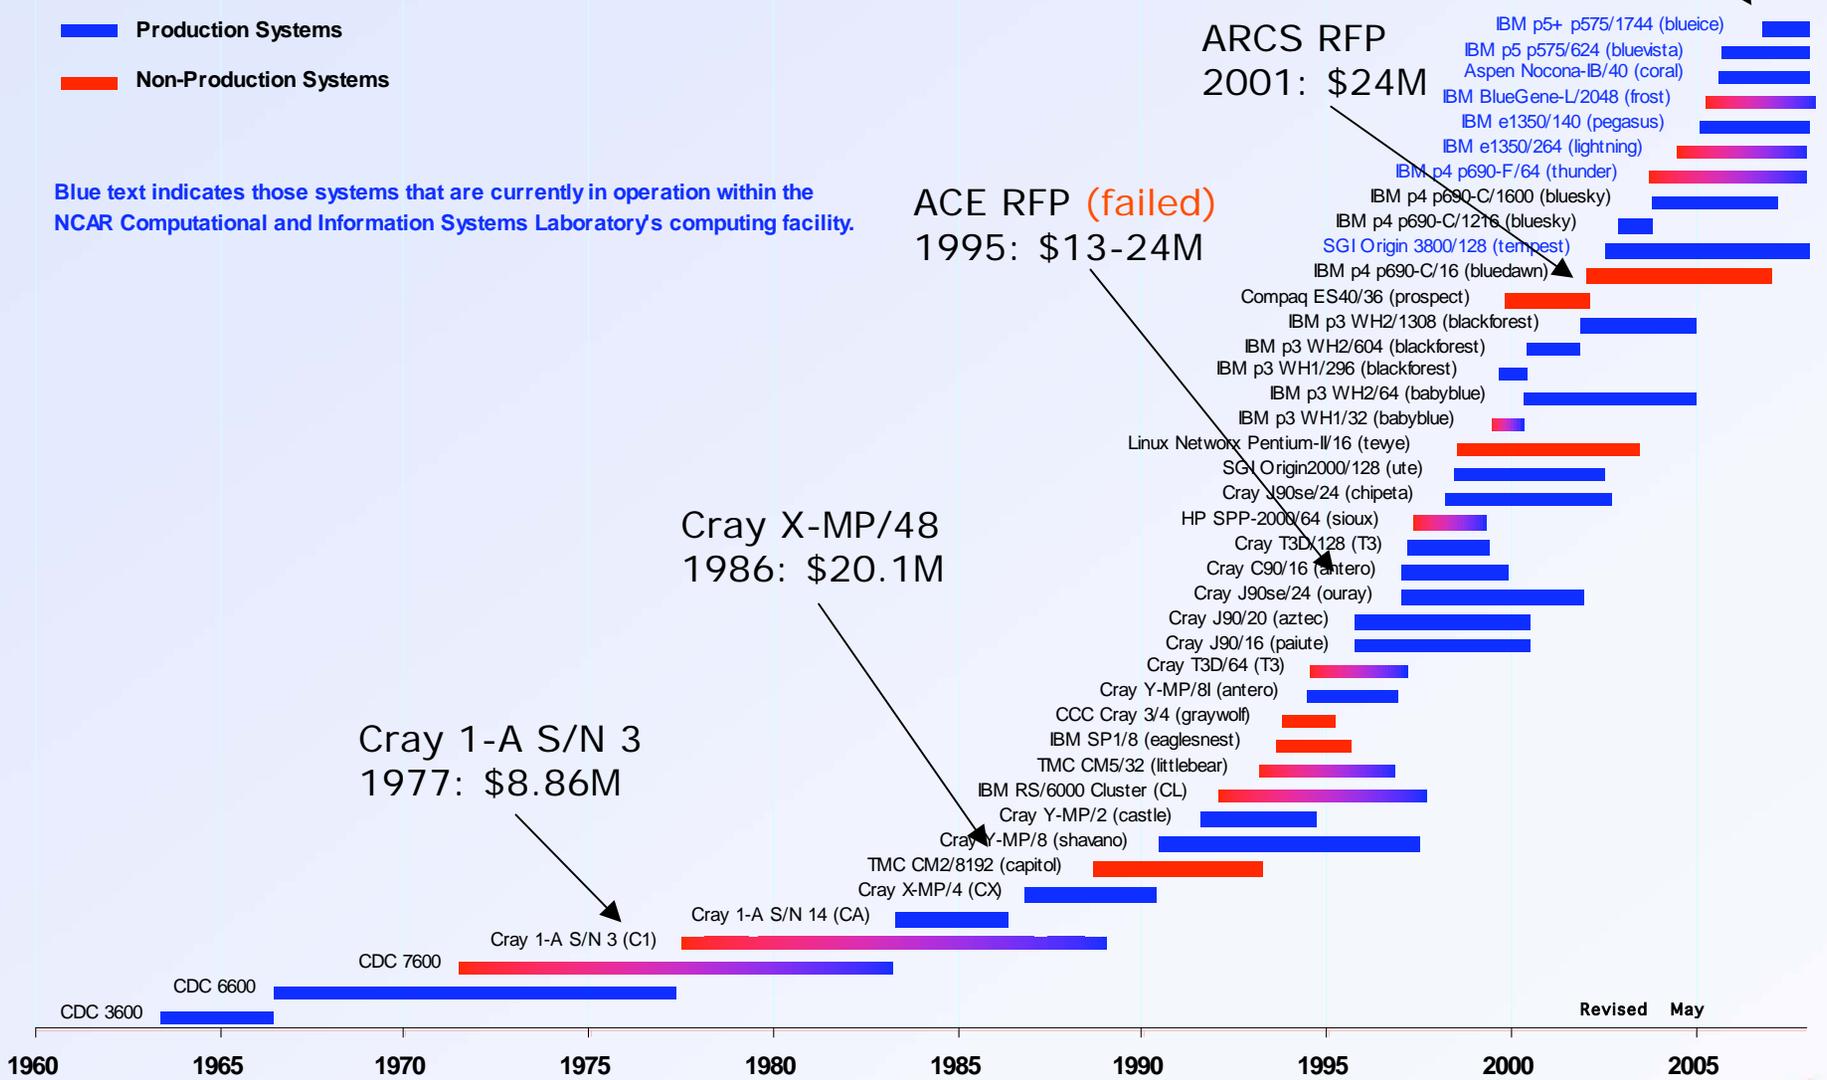
Tom Bettge
Director – Operations & Services Division
Computational and Information Systems Laboratory
National Center for Atmospheric Research

History of Supercomputing at NCAR

ICISS RFP
2006:
\$xxM

- █ Production Systems
- █ Non-Production Systems

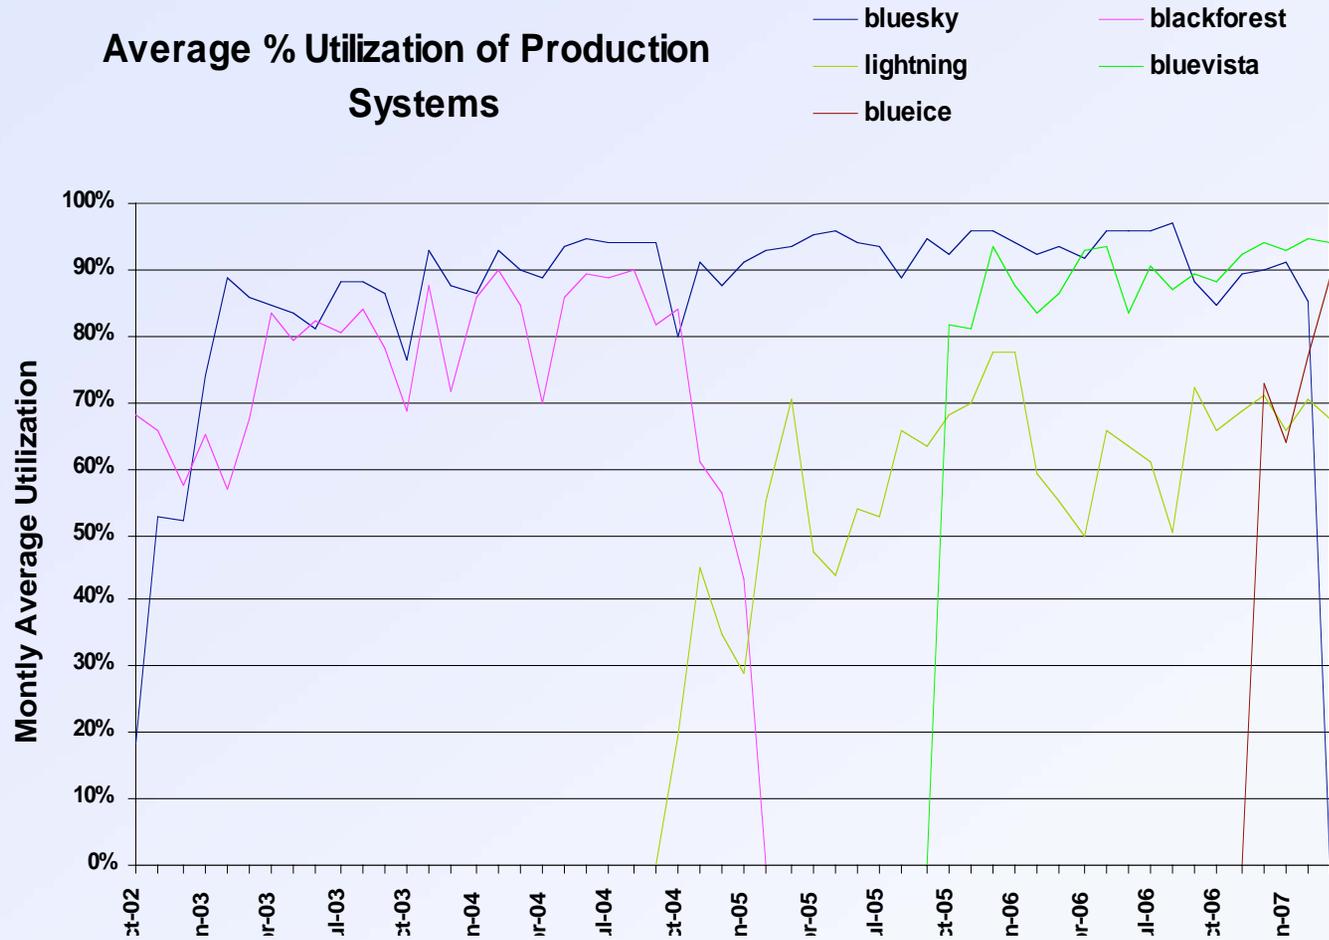
Blue text indicates those systems that are currently in operation within the NCAR Computational and Information Systems Laboratory's computing facility.



Servicing the Demand for HPC at NCAR

		Mar '07	2007 To Date	2006	2005
<ul style="list-style-type: none"> CISL's supercomputers are well utilized ... 	Blueice	89.1%	77.5%	-	-
	Bluevista	94.1%	94.0%	89.1%	-
	Lightning	67.3%	67.6%	63.3%	61.5%
	Bluesky 8-way LPARs	-	90.4%	91.7%	92.5%
	Bluesky 32-way LPARs	-	83.3%	92.9%	94.6%
<ul style="list-style-type: none"> ... yet average job queue-wait times are measured in hours and/or mins 		Regular Queue	March '07 Average Queue Wait Time	Lifetime Average Queue Wait Time	
	Blueice		1h48m	35m	
	Bluevista		2h04m	1h49m	
	Lightning		45m	11m	

Utilization by Machine



Outline

- NCAR Facility Specifications (Physical)
 - Power, cooling, space
 - Challenges
- ICESS Installation and Testing (Virtual)
 - Benchmarking
 - Installation (recently phased – or phased decommissioning)
 - Acceptance Testing
 - Integration into NCAR environment
 - Production Operations ('constant integration')
- Applications (Science)
- NCAR's Plans for New Facility

Facility History

- Original Computer Room – 1965
 - CDC 3600, 6600
- Major Expansion – 1975
 - Cray 1, XMP, YMP
- Significant Electrical Upgrade – 1983
- UPS Replacement -1997
- Generator Installation – 2004
- Chilled Water Plant Retrofit - 2005
- Return of Liquid Cooling - 2007



NCAR Facility Specs

Size

- ~7000 sq. ft. 36" raised floor
- ~6000 sq. ft. 17" raised floor
- ~2000 sq. ft. slab (Tape Archive Systems)

Power

- **1.2MW UPS capacity (Computer Load)**
- 2.5MW Total (Computing & Cooling)

Cooling

- 550 Ton cooling tower – dual cells for redundancy
- 480 Ton chilled water plant
- Mixture of IBM Coolblue Doors & Typical Liebert AHU

Efficiency

- $PUE = \text{Total Power Consumption} / \text{IT Equipment Power Consumption}$
- Typically at 1.65 – 1.7
- More simply for every 100kW of computer load 65kW of cooling load

Facility Challenges

- Operating a facility near it's maximum design point
 - Capacity planning for new installations
 - Reserving capacity (avoid power creep!)
- Utility stability
 - 2 significant outages per year on generator for 2 hours or more
 - Usually snow or wind
 - 1-2 momentary outages (too short to auto-start generators)
 - Can lose chilled water plant
 - Non-operational temperatures in < 7 minutes
- Return of liquid cooling
 - Fortunately had in-house expertise still available
 - Quality installation a must
- Trust but verify
 - New equipment installations
 - Verify manufacturers torque specs – even factory installed breakers
 - Generator control packages
 - verify, verify, verify sequences of operation

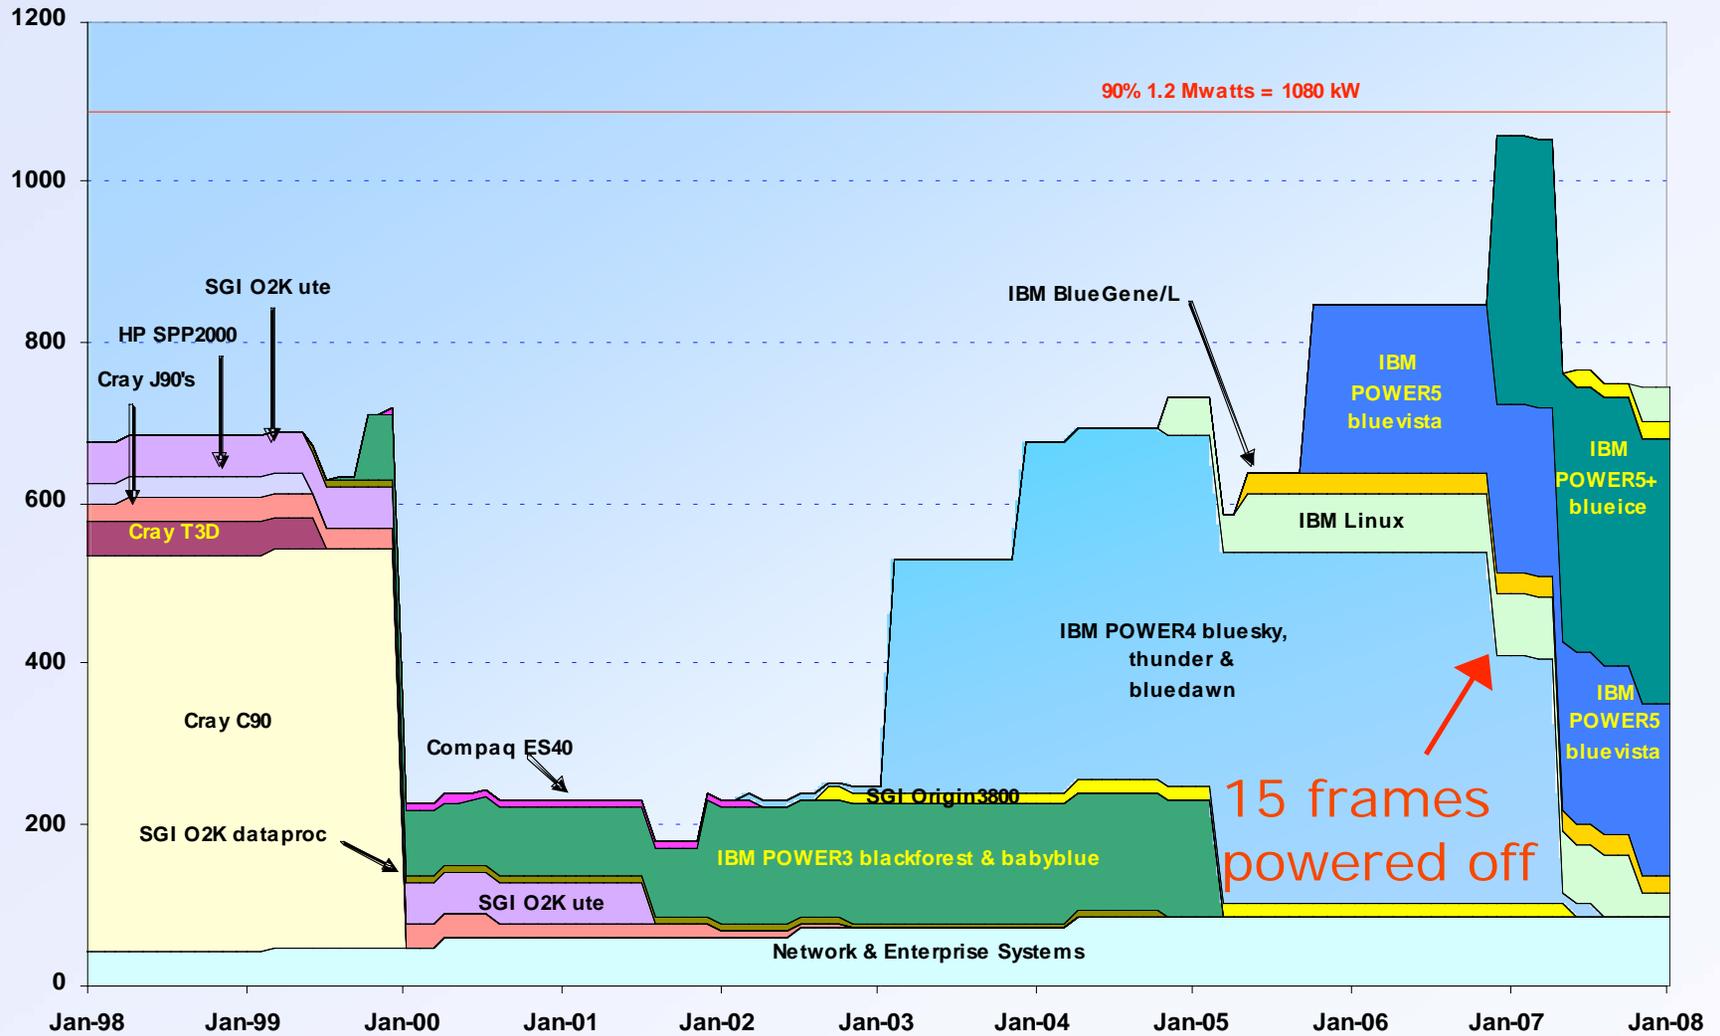
ICESS Installation and Acceptance Testing

October 2006 – August 2008

- Two computer systems (POWER5+ and POWER6) being delivered under the ICESS subcontract with IBM, but
- Three Acceptance Test Periods (ATPs):
 - One POWER5+ (blueice) ATP
 - ‘early decommissioning’ of fifteen frames of POWER4 (bluesky) to provide electrical ‘headroom’
 - Two POWER6 ATPs (installed in stages due to the Mesa Lab facility’s power & cooling limitations):
 - Stage 1: at least enough of the POWER6 system to run the POWER5+ system’s workload (more if the facility’s power & cooling limits will allow it) will be brought up and run through acceptance; the POWER5+ system’s workload then migrates to the first part of the POWER6 system
 - Stage 2: the POWER5+ system is powered off, and the remaining part of the POWER6 system is brought up and run through acceptance

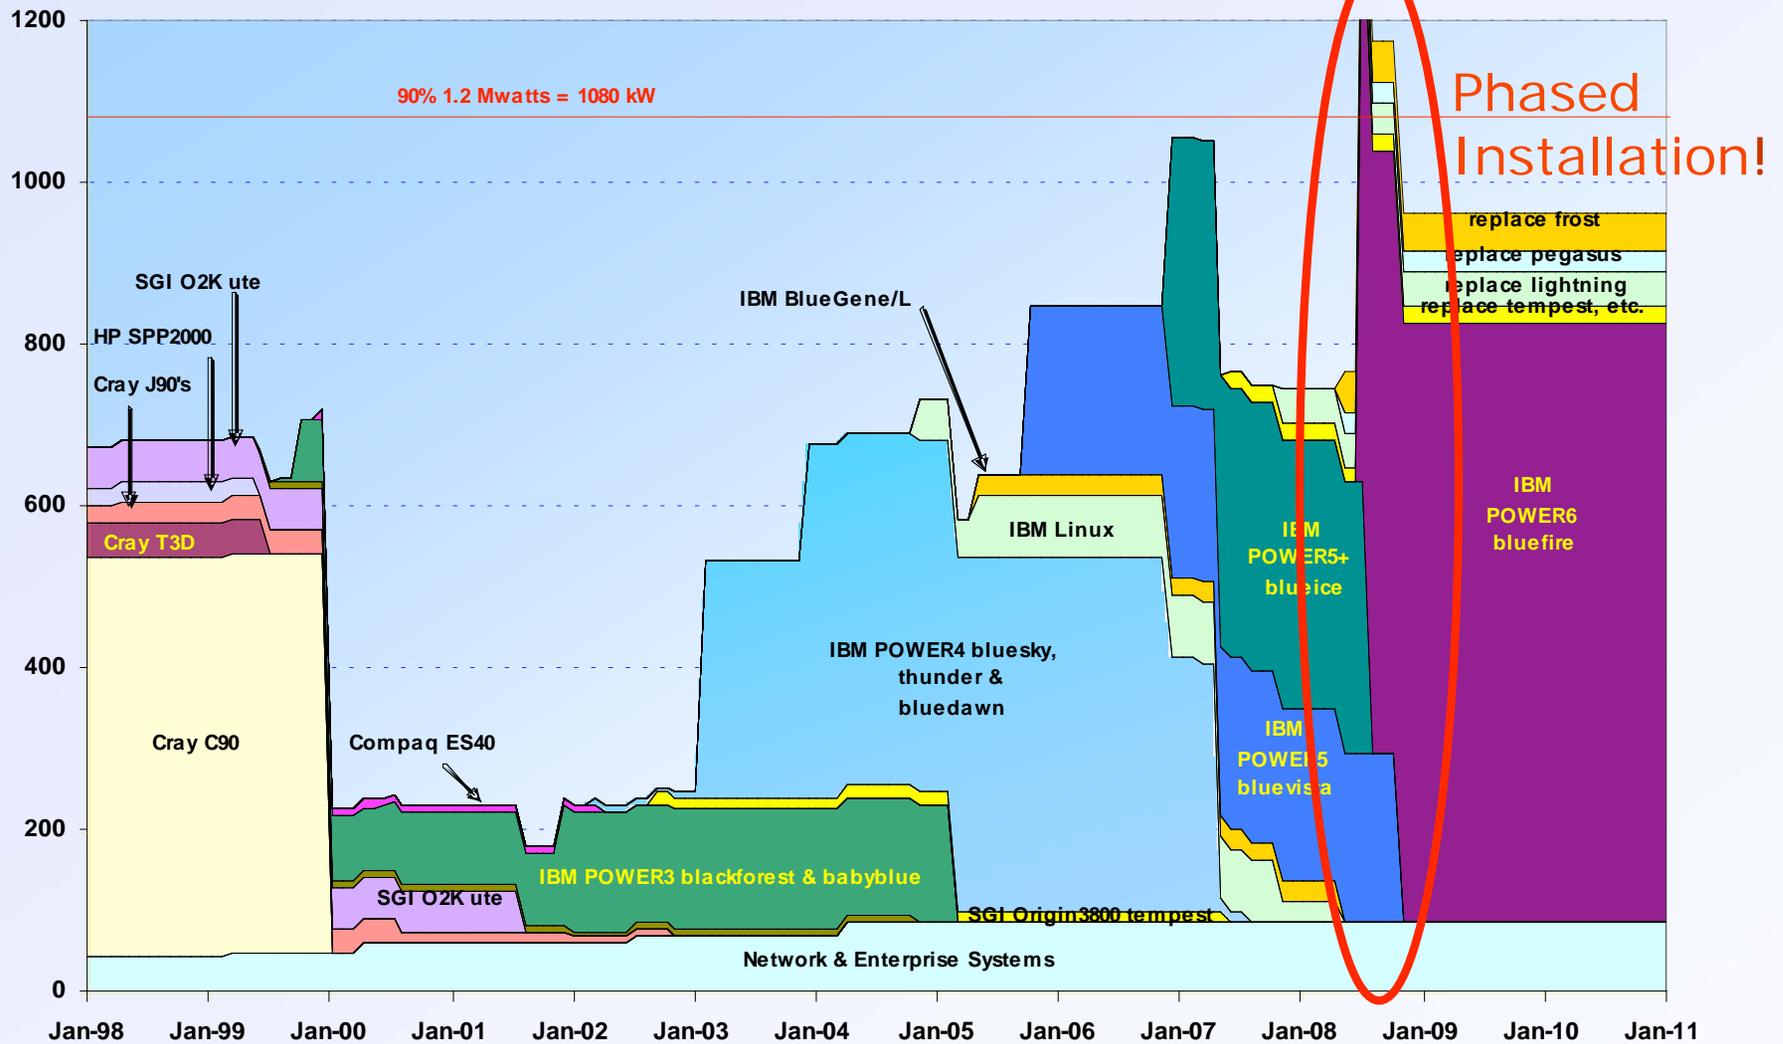
Power Profile

CISL Computer Facility Power Consumption (Measured kW)



Anticipated Power Profile

CISL Computer Facility Power Consumption (Measured kW)



System Testing (ICESS)

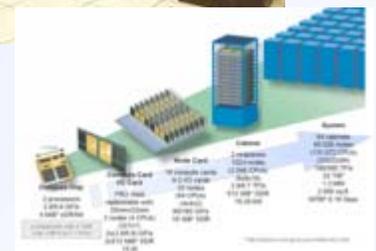
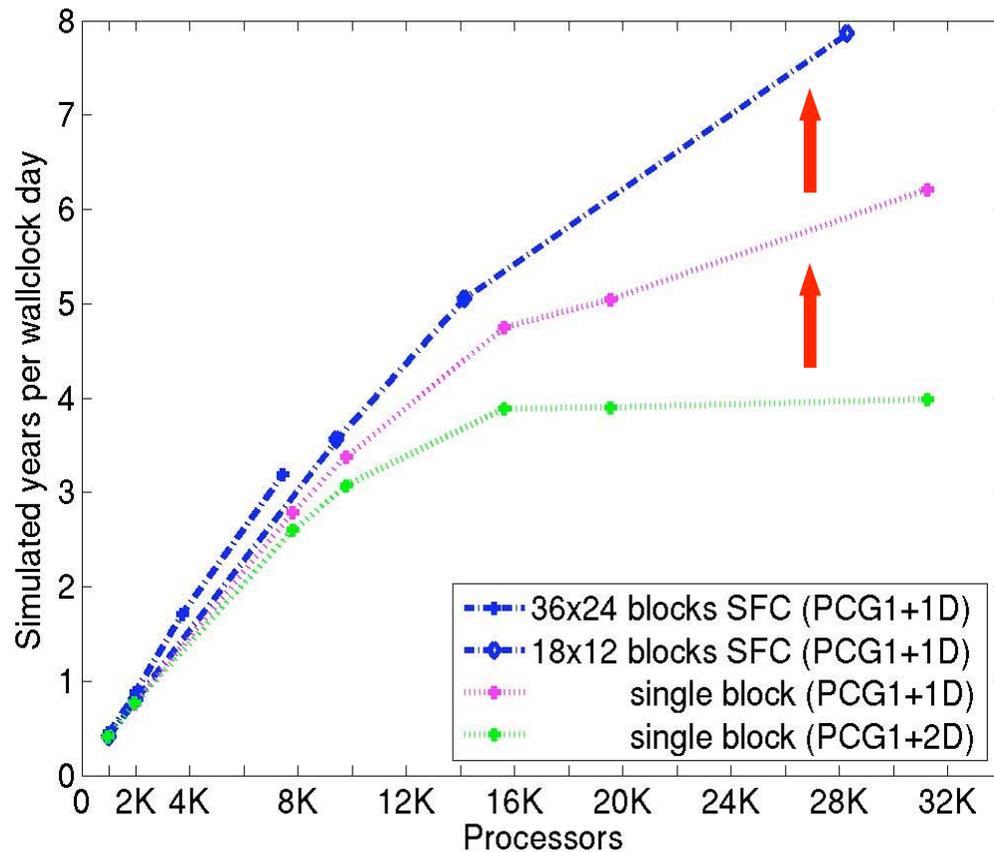
- Installation & Acceptance Testing
 - Physical Installation
 - *Phased power-up & monitoring of facility environmentals*
 - electrical & cooling loads
 - room air temperature & humidity
 - *over-temp alarms on neighboring systems!*
 - rebalance air handlers & airflow
 - Benchmark rerun by Vendor
 - *numerical correctness*
 - *performance must meet/exceed Vendor's promises*
 - capacity commitment (ICESS P5+ / P6) (benchmarks)
 - capability commitment (ICESS P5+ / P6) (benchmarks)
 - Acceptance Testing (P5+ / P6)
 - *Benchmark rerun by CISL (NCAR must be able to reproduce Vendor's results)*
 - *Workload Stress Test*
 - Duration: 14 continuous days
 - Availability > 98%, MTBF > 72 hrs

System Integration (ICESS)

- Configuration of User Environment and NCAR customization of Software Stack
- Grant Early Users Access
- Early Production
 - *mid January*
 - *early users and initial breakthrough science code testing*
- Capability Science Projects
 - *3/4 of resource dedicated for first 4 months of production*
- Full Production
 - *June 1st target*
 - *CSL, NCAR labs, Universities*
- Integration – through contract lifetime
 - Periodic BM reruns to detect regressions, degradations
 - Monitoring of service delivery
 - System Reliability must exceed negotiated criteria
 - *Availability > 98.5%*
 - *MTBF > 288 hrs*

Application Design.....

POP 0.1 Degree BG/L Improvement



John Dennis, NCAR

NCAR Advanced Computing Facility



NCAR



UCAR

Computational and Information Systems Laboratory - NCAR
Copyright © 2007 - University Corporation for Atmospheric Research

15 May 2007



Challenge Beyond 2007



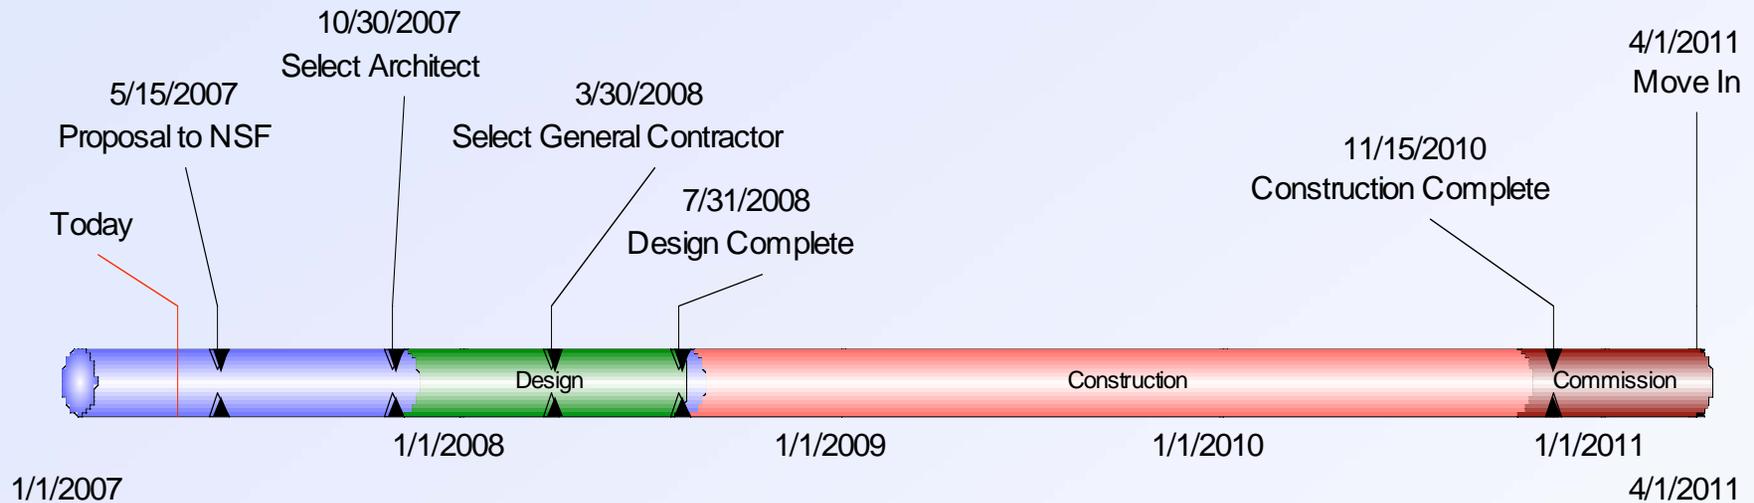
- **NCAR Computing Facility Expansion**
 - NCAR center limits of power/cooling/space will be reached with ICESS procurement addition – **2007-2008**
 - **2011** computing augmentation will **exceed center power capacity**
 - **Must plan and act now** (actually started planning in 2004)
- New center requirements have been established
- Conceptual design for new center has been completed
- **Partnership** – discussions with small number of potential university partners completed; funding options being finalized

NCAR Announces Data Center

TUESDAY JANUARY 23, 2007

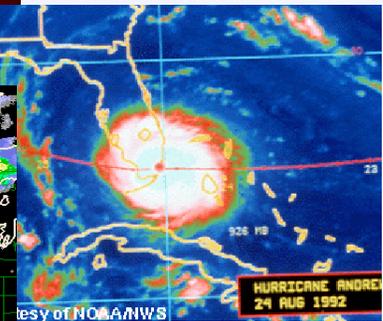
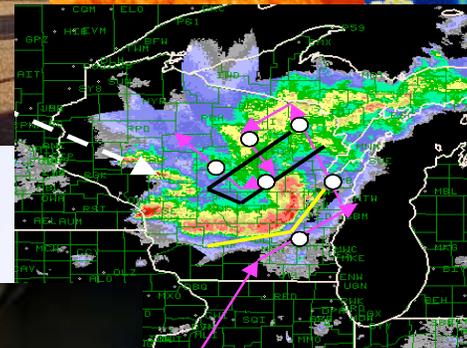
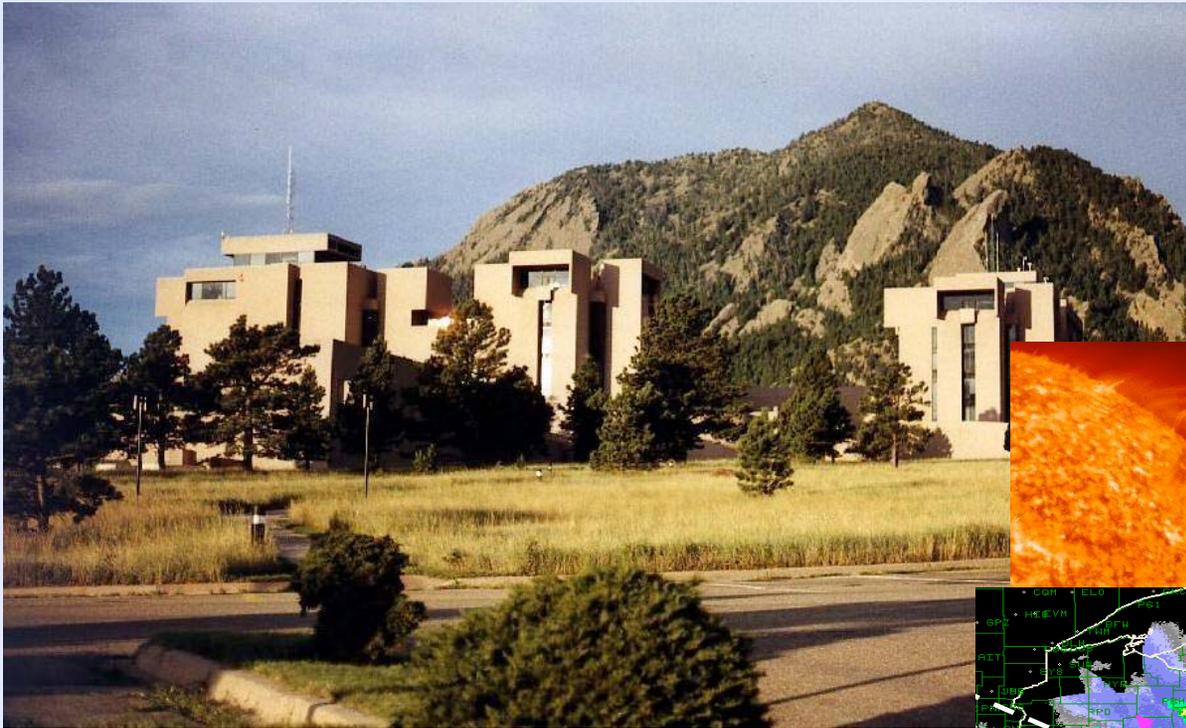
- NCAR announced that it has entered into a partnership with the [University of Wyoming](#) and the [State of Wyoming](#) to build a new supercomputing data center in Cheyenne.
- This new venture will replace the current end-of-life computer room with a modern facility that has [25,000 sq.ft. of raised floor and an initial capacity up to 8 megawatts of power](#). The new facility will ensure that NCAR's future is bright and productive in the provision of scientific cyber-infrastructure for research into climate, weather, and other geophysical scientific processes for the community.
- NCAR and Wyoming plan to build [new academic and scientific partnerships, as well as research universities nationwide](#). The data center's powerful infrastructure will enable research efforts to tackle important environmental problems.

Timeline



Targeted milestones:

Construction initiation (“shovels in the ground”) – Late summer 2008
Facility open for operation – *Late 2010 / Middle 2011*



Thank You!

