


Deploying Large Scale Cray XT Systems at ORNL

Presented to the
Cray User Group

Jim Rogers
Director of Operations for the
National Center for Computational Sciences

Oak Ridge National Laboratory
Oak Ridge, Tennessee

May 7, 2009



The Jaguar Cray XT5 System



200 Cray XT5 Cabinets (25x8)

48 Liebert XDP Cooling Systems


37kW/cabinet – 7.5MW Total

184 ECOPhlex-cooled Cabinets

16 Air-cooled Cabinets

4500 square feet (system min. only)

The Kraken XT5 System



88 Cray XT5 Cabinets (22x4)

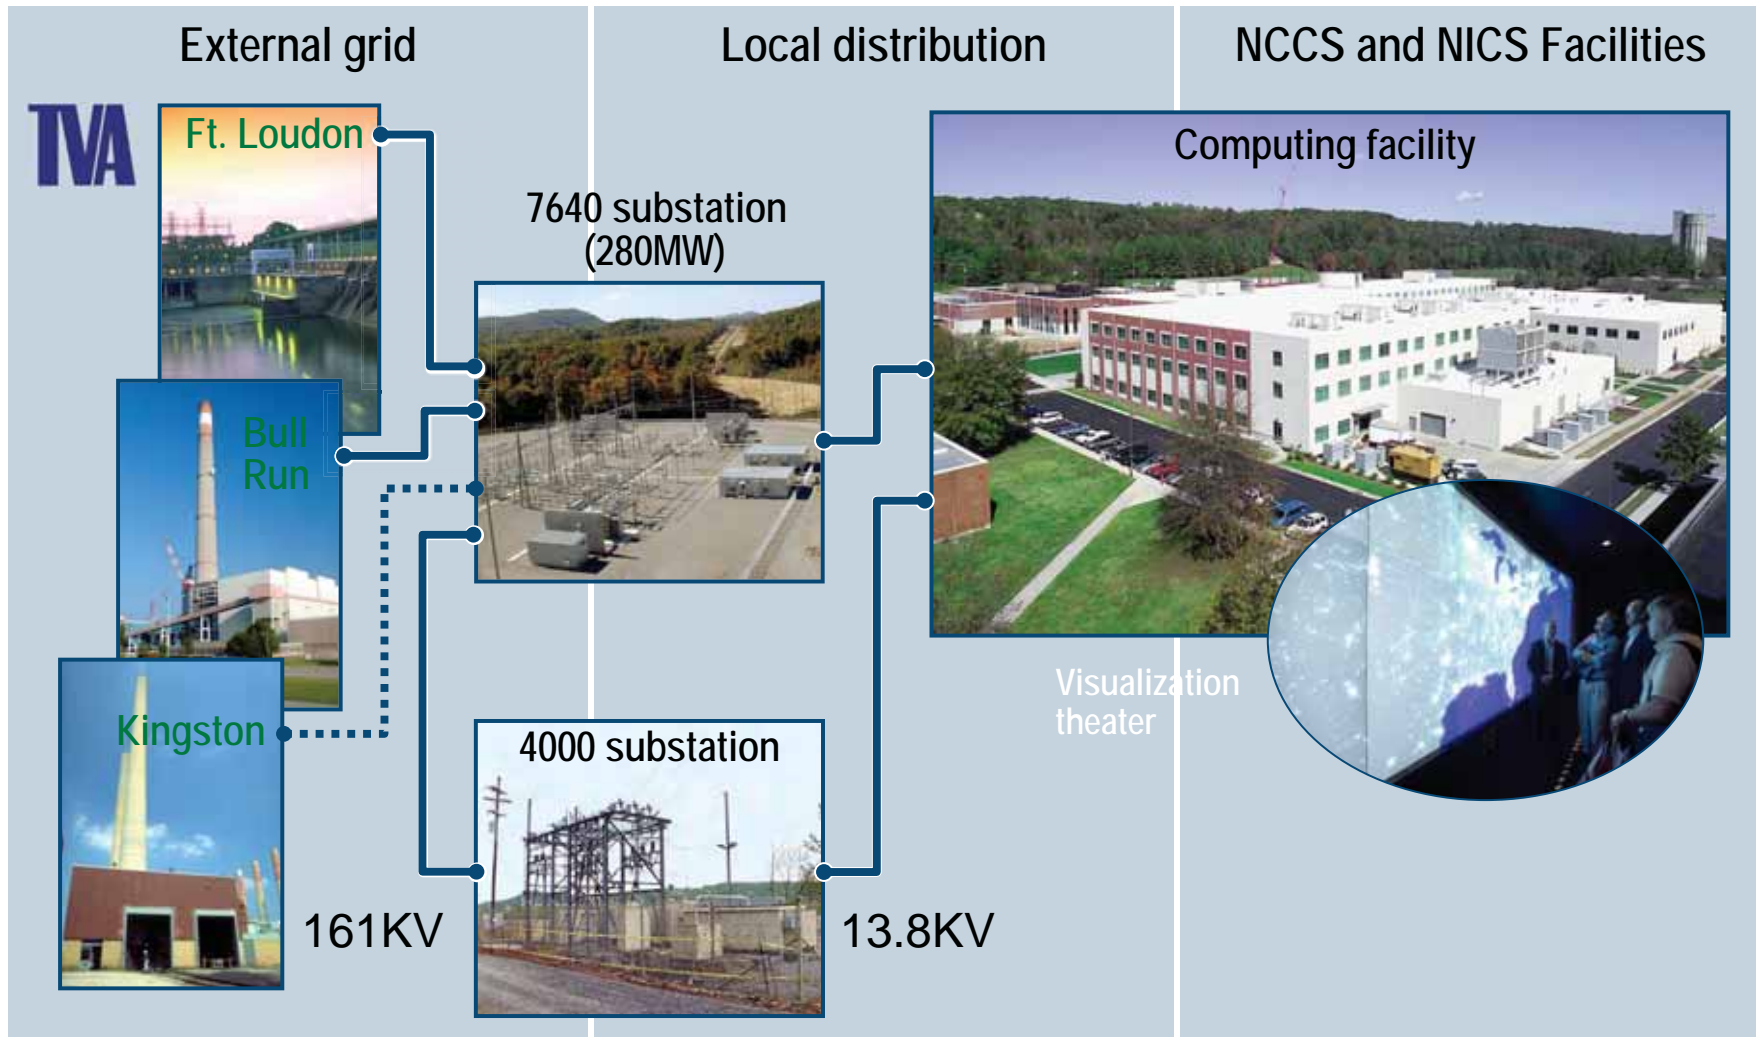
- 20 Liebert XDP Cooling Systems
- 37kW/cabinet – 3.3MW Total
- Complete ECOPhlex-cooled Solution
- 4-cabinet/XDP and 5-cabinet/XDP Configurations
- 1800 square feet (system min. only)

National Center for Computational Sciences - Preparing for two large Cray XT5 systems



- 40,000 ft² Aggregate Raised Floor Capacity
- 20MW Electrical Capacity (revised)
- 6,600 Tons of Chilled Water Capacity (revised)
- 1.5MW UPS Capacity (revised)
- 2.25MW Generator Capacity (revised)

ORNL Leverages TVA's Robust Electrical Infrastructure




Utility and Electrical Distribution Upgrades

- Upgrade ORNL 7640 electrical substation capacity from 210 to 280 MW
- Install two 13,800V distribution lines from 7640 to the NCCS
- Integrate new distribution lines into existing switch gear
- Harden the distribution infrastructure, eliminating the dependency on the 4000 substation

Transformers and Switchgear



Distribution equipment arrives as part of the 15MW expansion of electrical capacity.



Electricians work to position a 2500kVA transformer, one of two that will provide 480V service to the NICS Cray XT5.

CEP Upgrades to Chilled Water



Supplemental chilled water from an adjacent facility (typically 700-800 tons) offsets much of the existing comfort load, allowing the 5600 chillers to manage the heat load associated with the computer rooms



One of two new 1500 ton chillers is moved into the 5600 Central Energy Plant. Each new chiller has the capacity to remove more than 5.25MW of heat.

The two new 1500 ton chillers are operational. They typically run at 95-98% of capacity, with flow rates of ~6600gpm at 42° Fahrenheit. Δt is approximately 15° Fahrenheit

UPS and Generator Power



In 2HCY09, a series of CRUs and chilled water pumps will be placed on generator so that long term power outages will not disrupt the delivery of chilled water. Based on current heat load from UPS-protected systems, we anticipate no less than 30 minutes of additional time to complete administrative controls such as graceful shutdowns of file systems and infrastructure.

A 2200 HP, 1.5MW Diesel Generator is set on its pad outside 5600. Aggregate diesel generator capacity is 2.25MW

9 A 1.0MW Rotary UPS is installed in the primary electrical distribution room inside 5600. Aggregate UPS capacity is 1.5MW, protecting disks and critical infrastructure



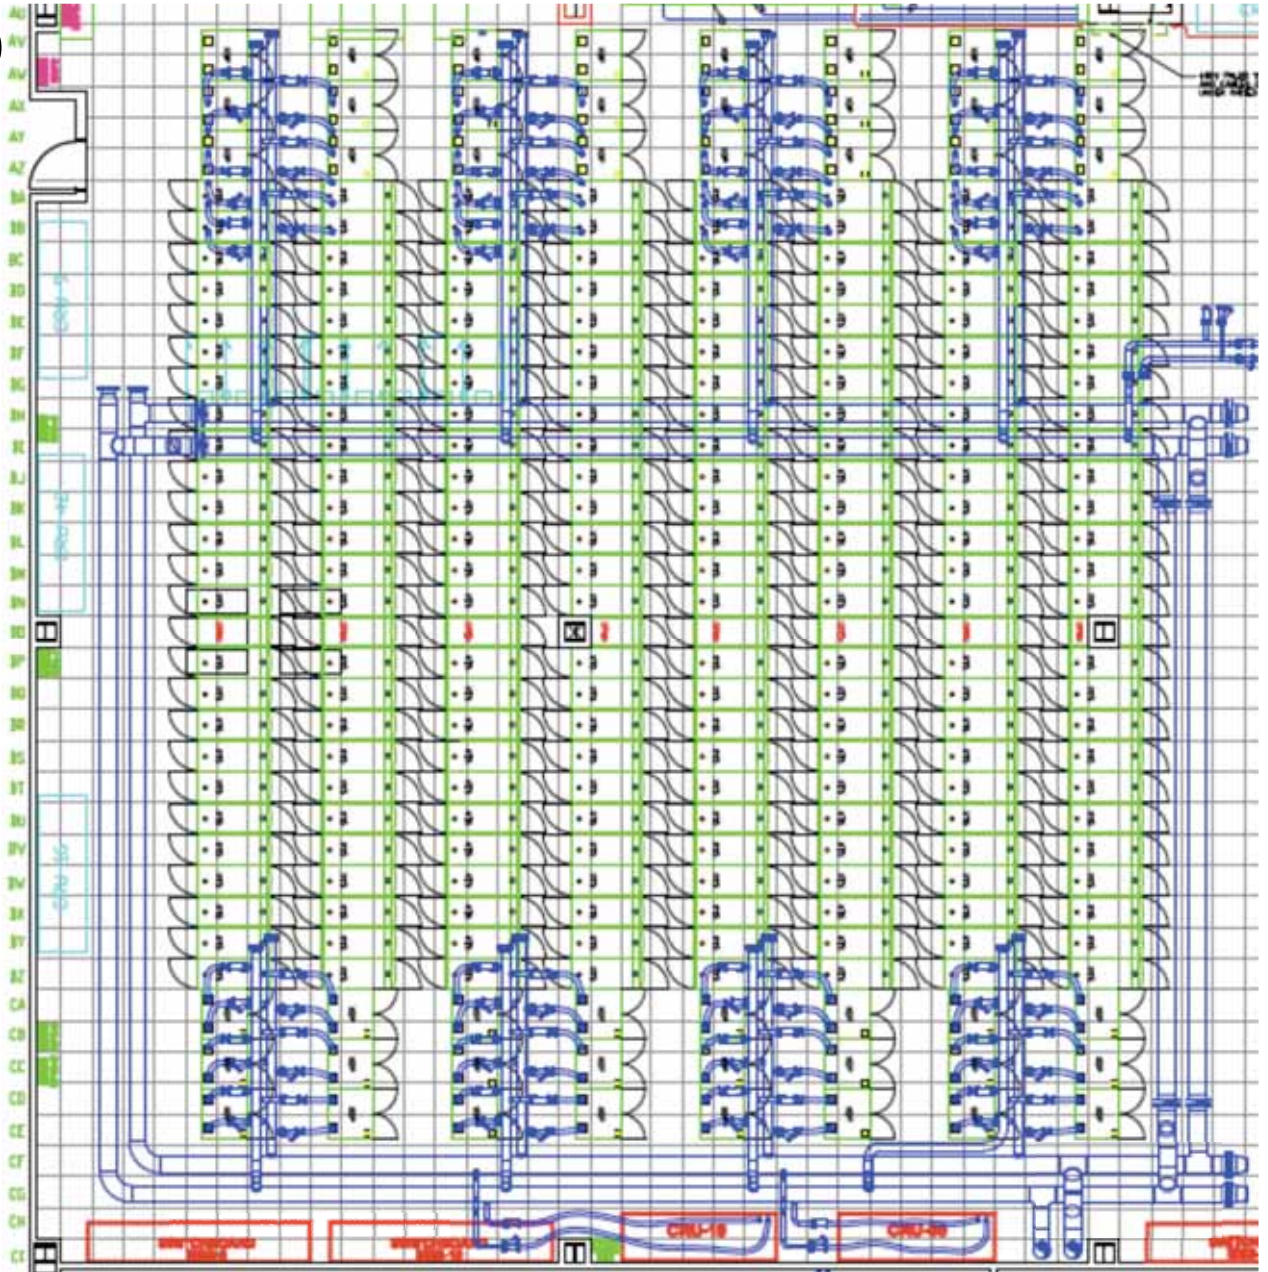
Tying Transformers to the Switchboards



Electricians complete the connections from the 2500kVA transformers to the three switchboards necessary for Jaguar. The completed 12" chilled water supply and return lines, with insulation installed are shown bottom right. The transformers are located immediately behind the switchboards, substantially reducing the run length of the feeds, saving installation dollars and operating costs.

Mechanical Upgrades for Jaguar

- Replaced more than 200 linear feet of 10" CW Line with new 12" CW Loop (44% increase in volume)
- Added Valves to Allow Segmenting of the Room and Prevent Subsequent CW Outages for New Systems
- Added 8 header systems for the 240kW Liebert XDPs
- Reoriented chilled water connections for multiple CRUs



Mechanical Upgrades for Jaguar




Workers use a small lift to remove old 10" chilled water piping, and replace it with 12" piping beneath the 36" raised computer room floor.

10" chilled water pipe is prepped for removal. A new 12" supply line will feed the first floor computer room, supplementing the original 12" line.



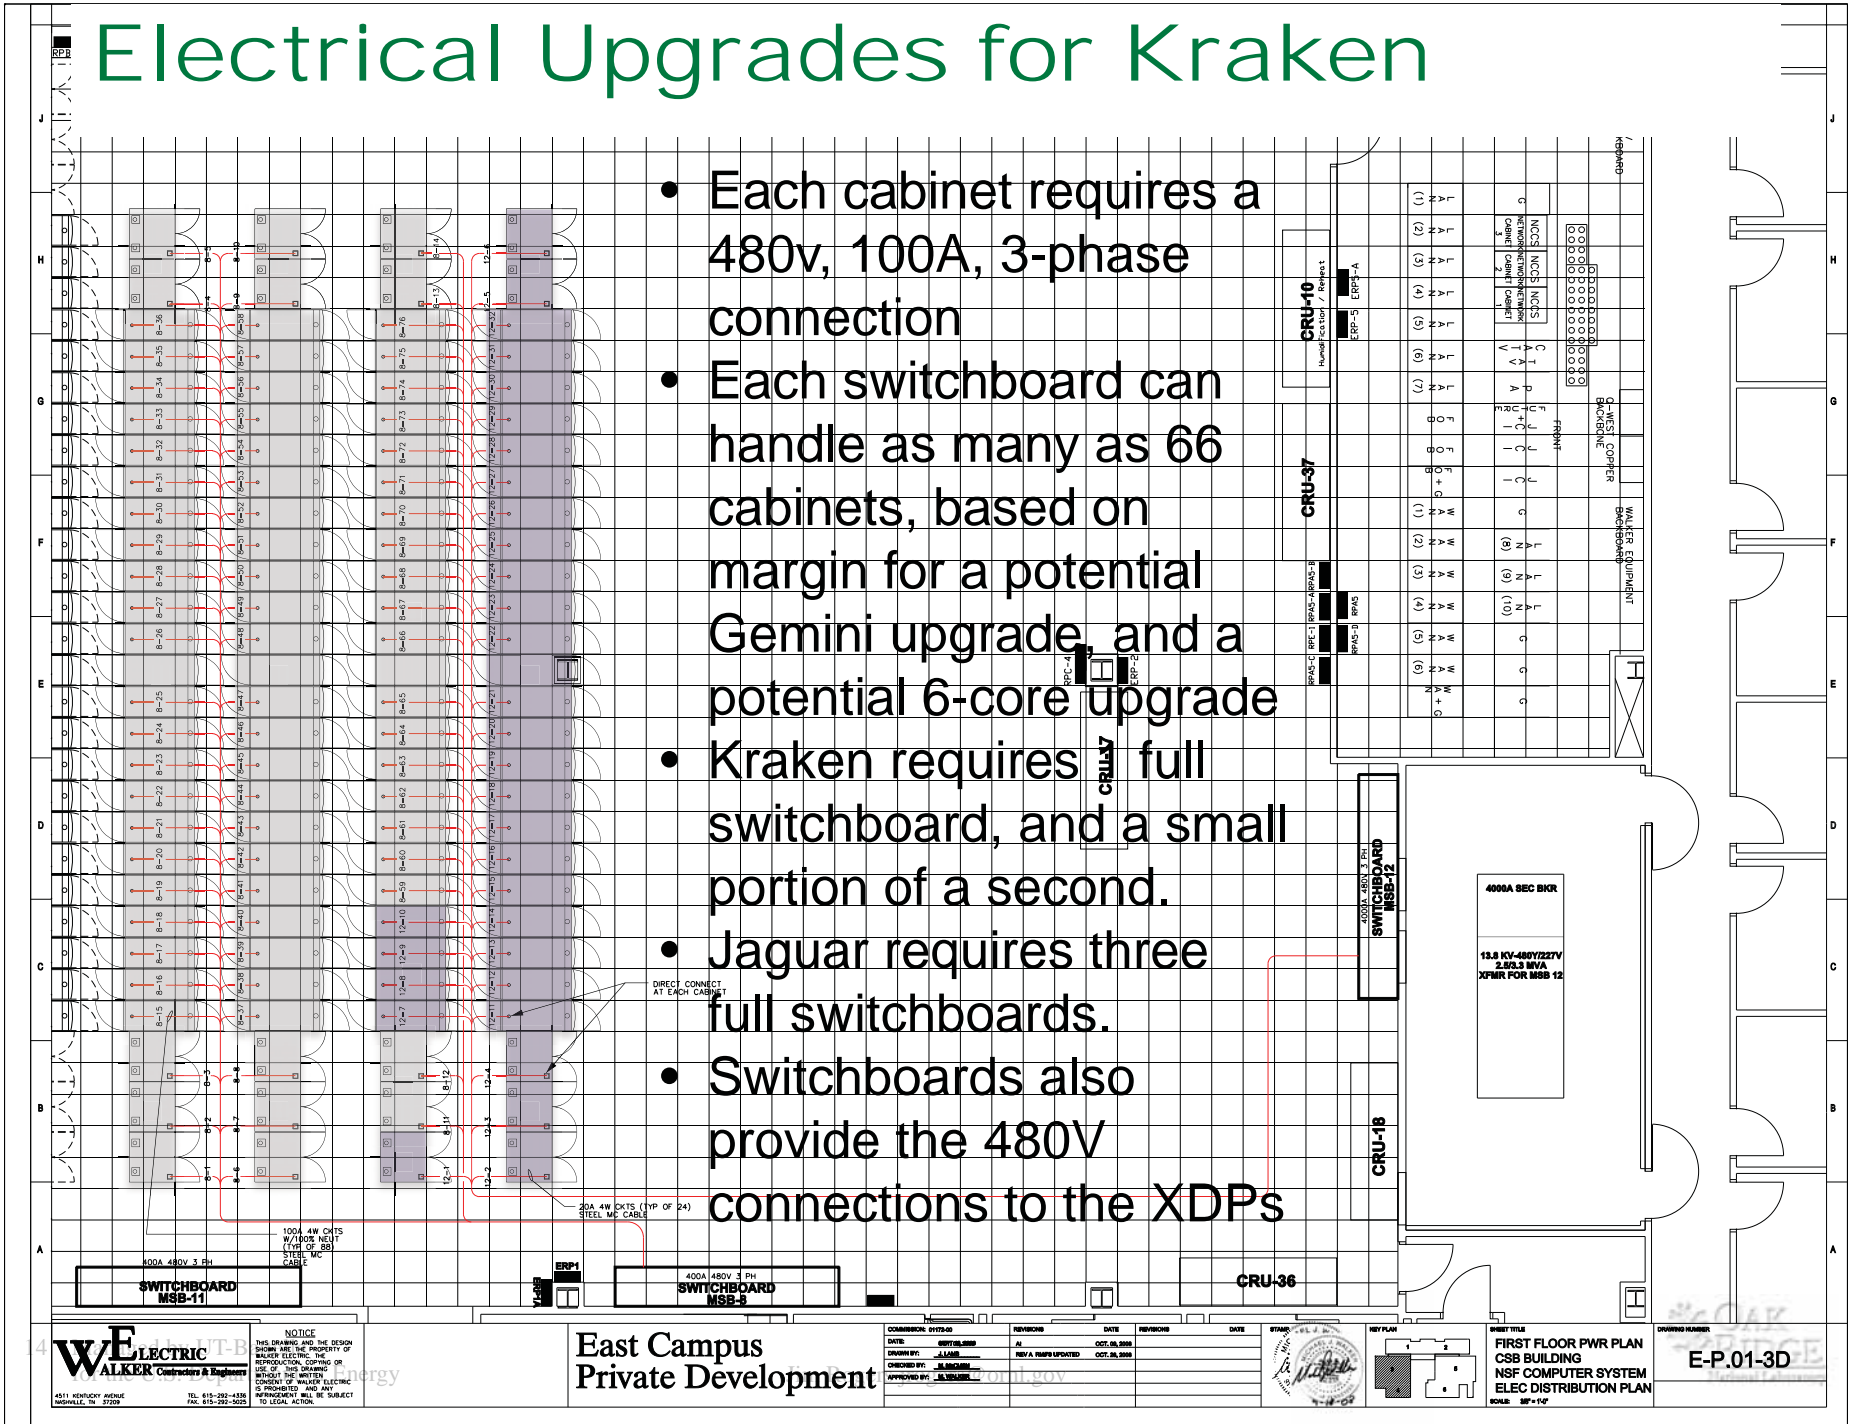
Mechanical Upgrades for Kraken

- Kraken, at 88 cabinets and 20 XDPs, did not require the replacement of the existing 10" lines, but did force the decommissioning of the Cray X1E, and its chilled water piping
- Installed six more stub headers to support the 240kW Libert XDPs. This configuration provided the flexibility for Kraken to be installed adjacent to Jaguar, or with up to about twelve feet of separation.



Electrical Upgrades for Kraken

- Each cabinet requires a 480v, 100A, 3-phase connection
- Each switchboard can handle as many as 66 cabinets, based on margin for a potential Gemini upgrade, and a potential 6-core upgrade
- Kraken requires full switchboard, and a small portion of a second.
- Jaguar requires three full switchboards.
- Switchboards also provide the 480V connections to the XDPs




WALKER ELECTRIC
 WALKER Contractors & Engineers
 4511 KENTUCKY AVENUE
 INDIANAPOLIS, IN 46209
 TEL: 317-252-4226
 FAX: 317-252-2623

NOTICE
 THIS DRAWING AND THE DESIGN
 THEREON ARE THE PROPERTY OF
 WALKER ELECTRIC. NO REPRODUCTION,
 REPRODUCTION, COPYING OR
 USE OF THIS DRAWING
 WITHOUT THE WRITTEN
 CONSENT OF WALKER ELECTRIC
 CONTRACTORS & ENGINEERS
 IS PERMITTED. ANY
 REPRODUCTION WILL BE SUBJECT
 TO LEGAL ACTION.

East Campus
 Private Development

COMMISSION	DATE	REVISIONS	DATE
0172-00			
DATE	BY	DATE	BY
DRAWN BY: J. LAMM	REV A REVISIONS	OCT. 04, 2009	
CHECKED BY: J. BLOOMER		OCT. 04, 2009	
APPROVED BY: J. BLOOMER			




SHEET TITLE
 FIRST FLOOR PWR PLAN
 CSB BUILDING
 NSF COMPUTER SYSTEM
 ELEC DISTRIBUTION PLAN
 30x36 3/4" x 1/4"

DRAWING NUMBER
E-P.01-3D

Switchboard Installation


Individual connections for the XT5 cabinets are prepped for connection to the switchboard prior to being run to final installation locations.



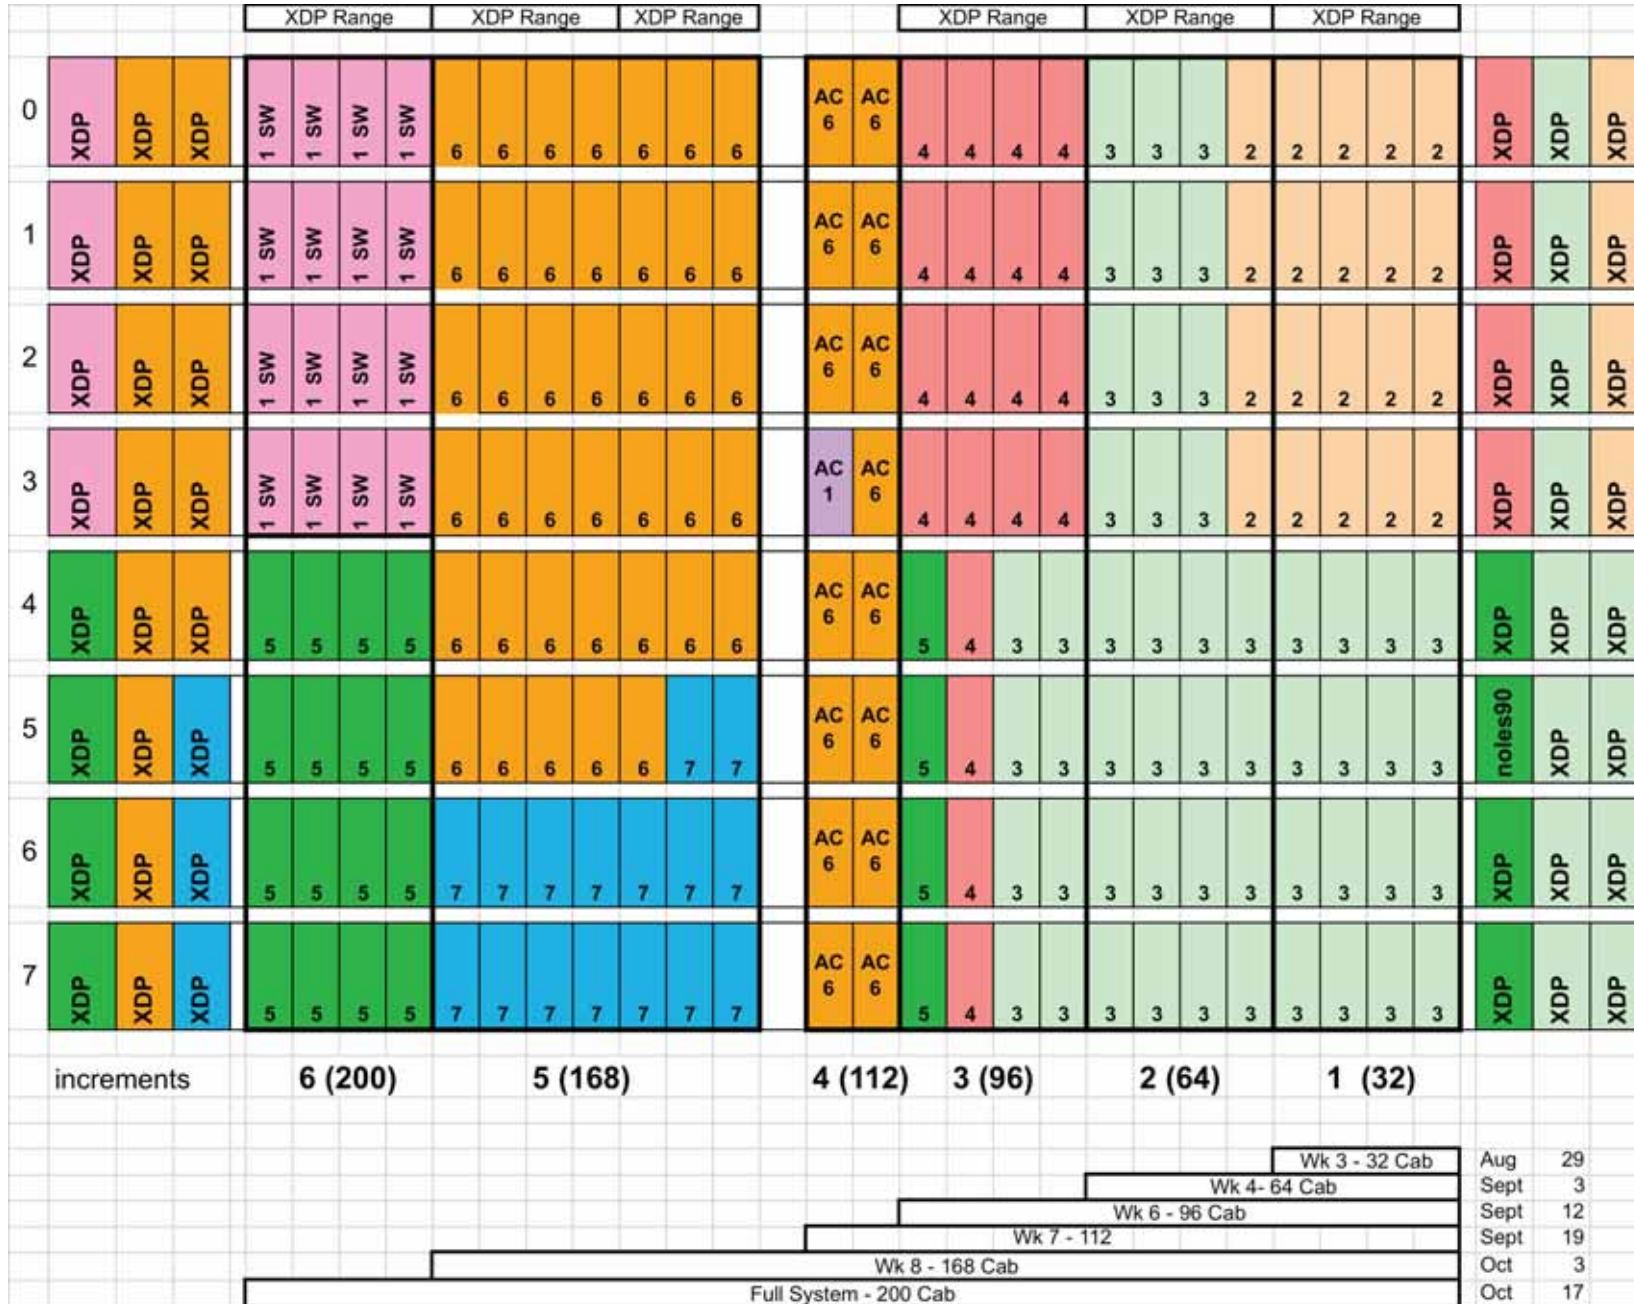
Fully dressed connections from a switchboard back to the transformer

Modeling Room Temperatures with TileFlow

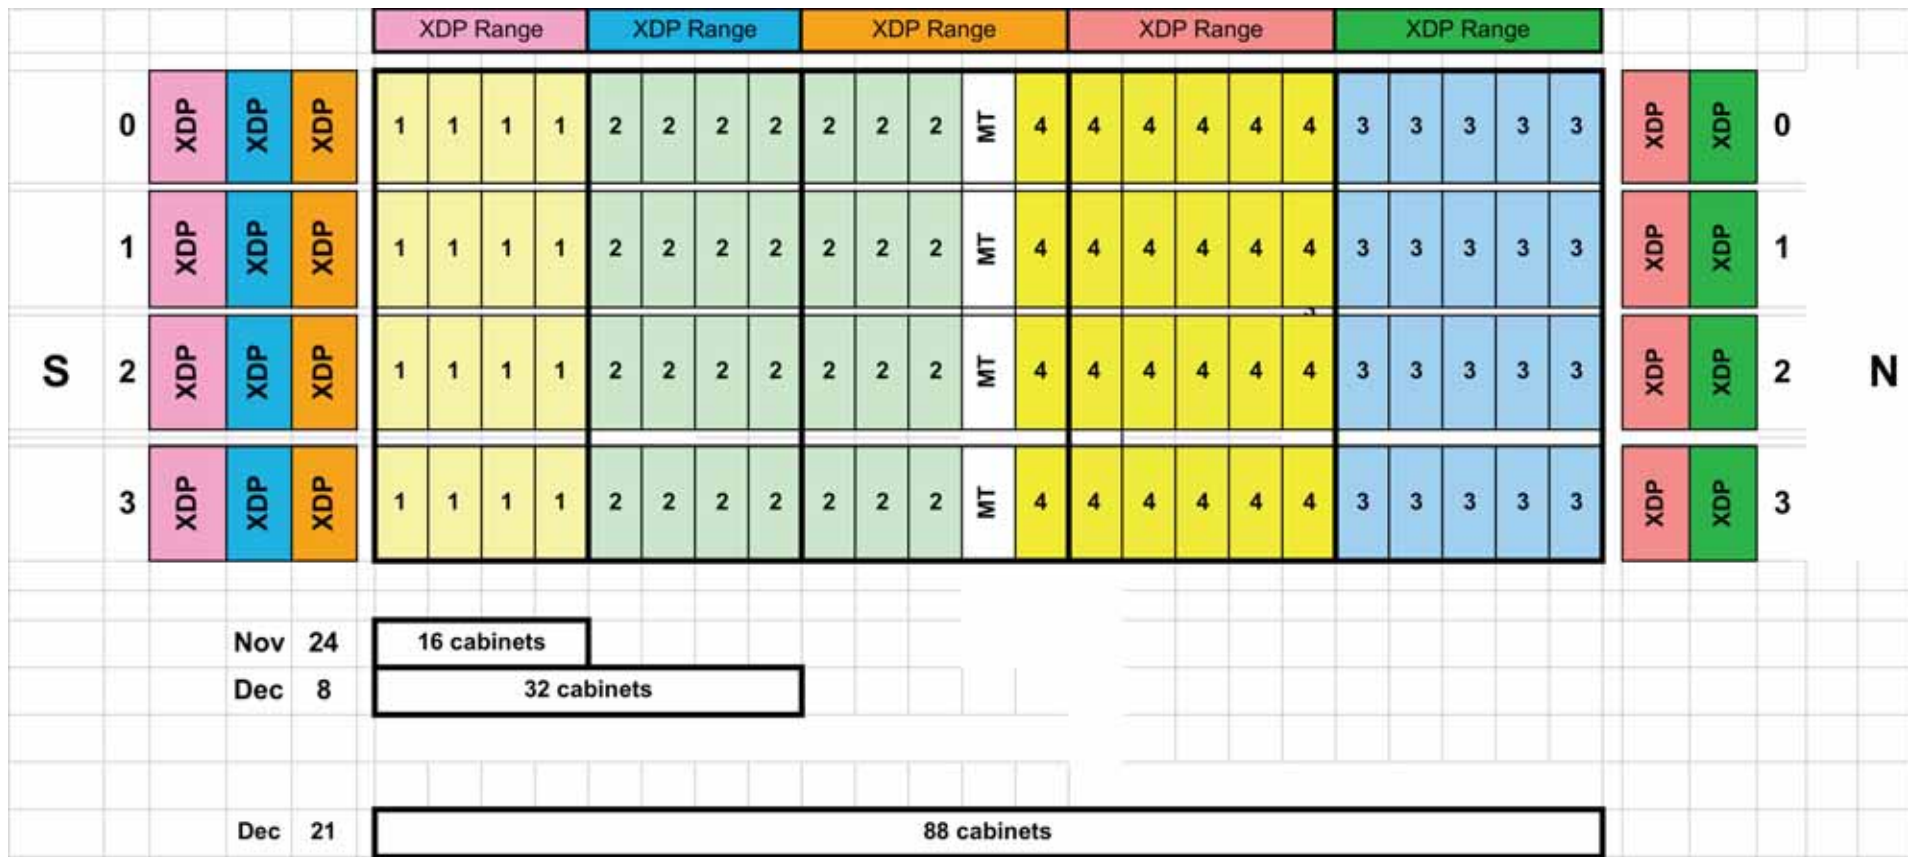
- Jaguar is a split system
 - 184 ECOphlex™-cooled cabinets
 - 16 air-cooled cabinets
 - ORNL modeled below-floor air temperature and static pressure, and affect on inlet temperature based on placement of perforated tiles
- Kraken is fully ECOphlex™-cooled
- TileFlow identified needs for rearranging CRUs to handle the 10% ejected heat



Jaguar XT5 Installation Plan



Kraken XT5 Installation Plan

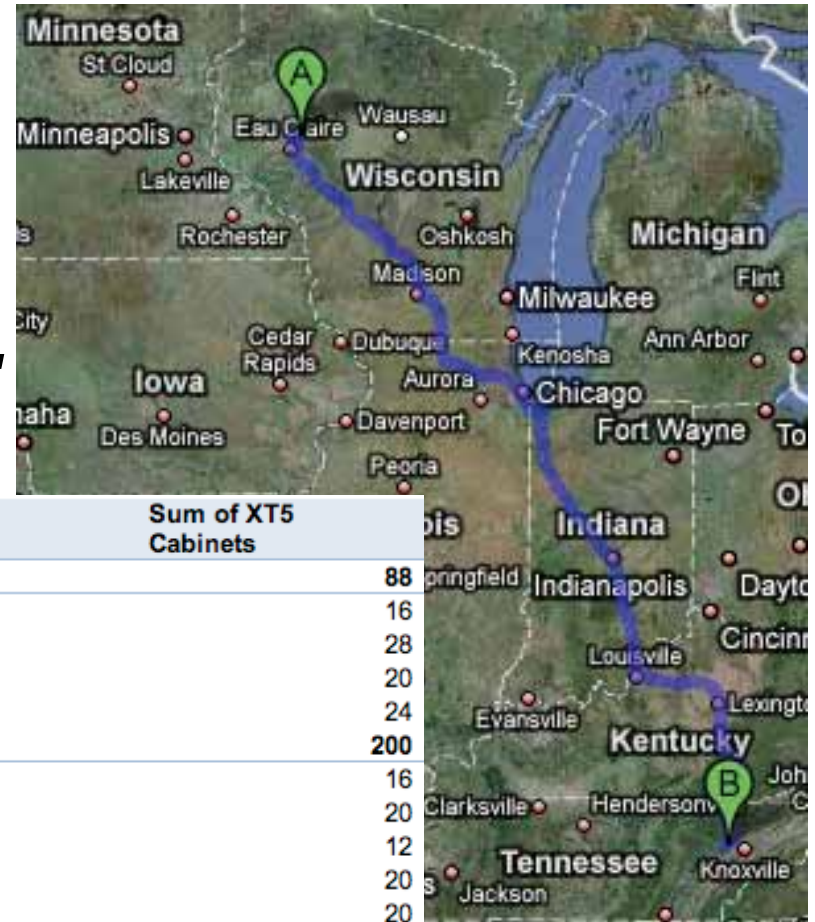


Deliveries

- Fun with OTR

- Driving directions from Cray Inc., Chippewa Falls, WI to 1 Bethel Valley Rd, Oak Ridge, TN 37830

- 866 mi - about 14 hours 3 mins
- 290 cabinets
- 28 separate deliveries
- 48,000 miles
- Does not include the Liebert XDP deliveries from Columbus, OH.



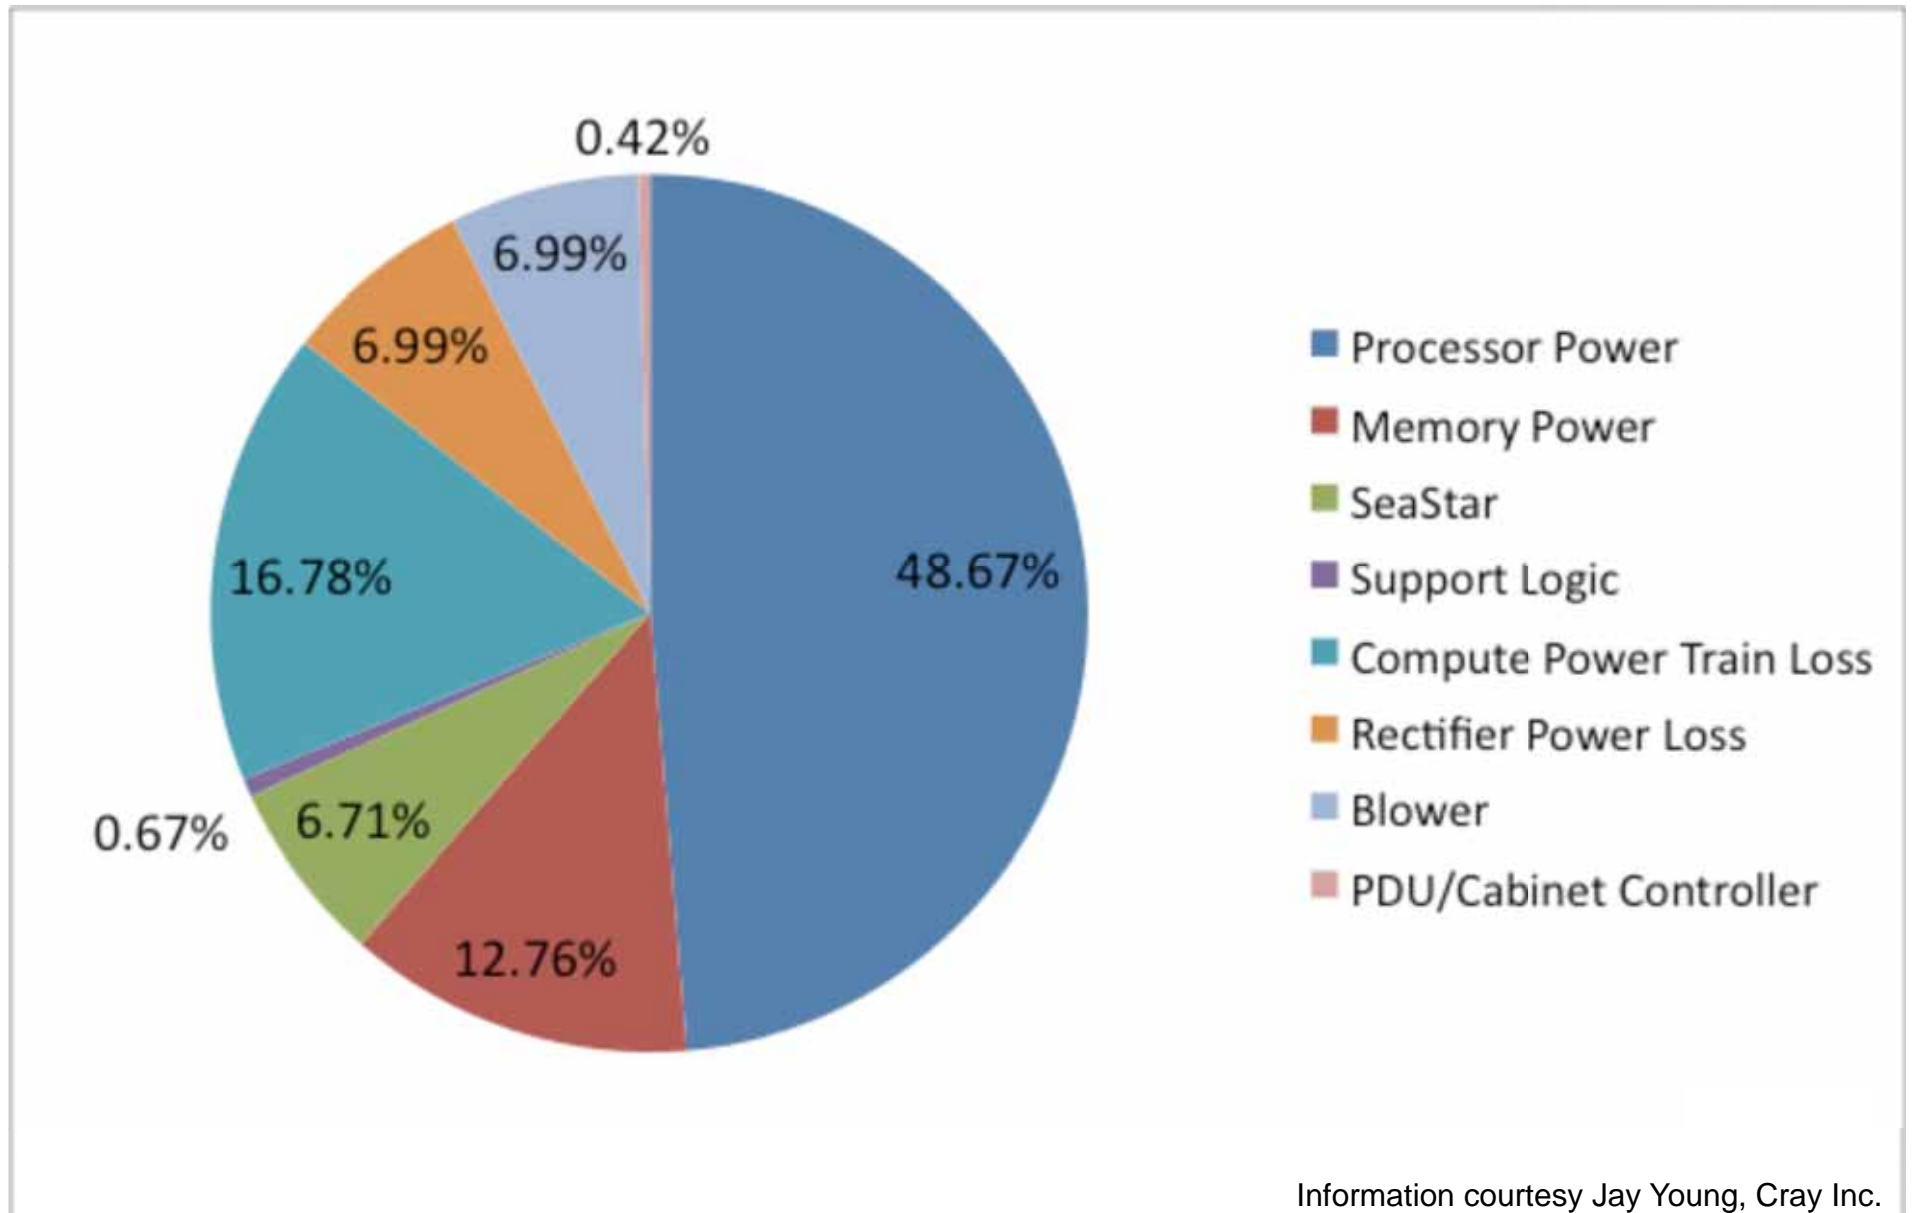
Row Labels	Sum of XT5 Cabinets
NSF	88
11/14/2008	16
12/1/2008	28
12/2/2008	20
12/12/2008	24
ORNL	200
7/29/2008	16
8/8/2008	20
8/20/2008	12
8/21/2008	20
8/22/2008	20
8/29/2008	20
9/4/2008	16
9/5/2008	20
9/10/2008	20
9/12/2008	20
9/16/2008	16
ORNL TDS	1
7/7/2008	1
ORNL XT6 - Hot Spare	1
7/29/2008	1
Grand Total	290



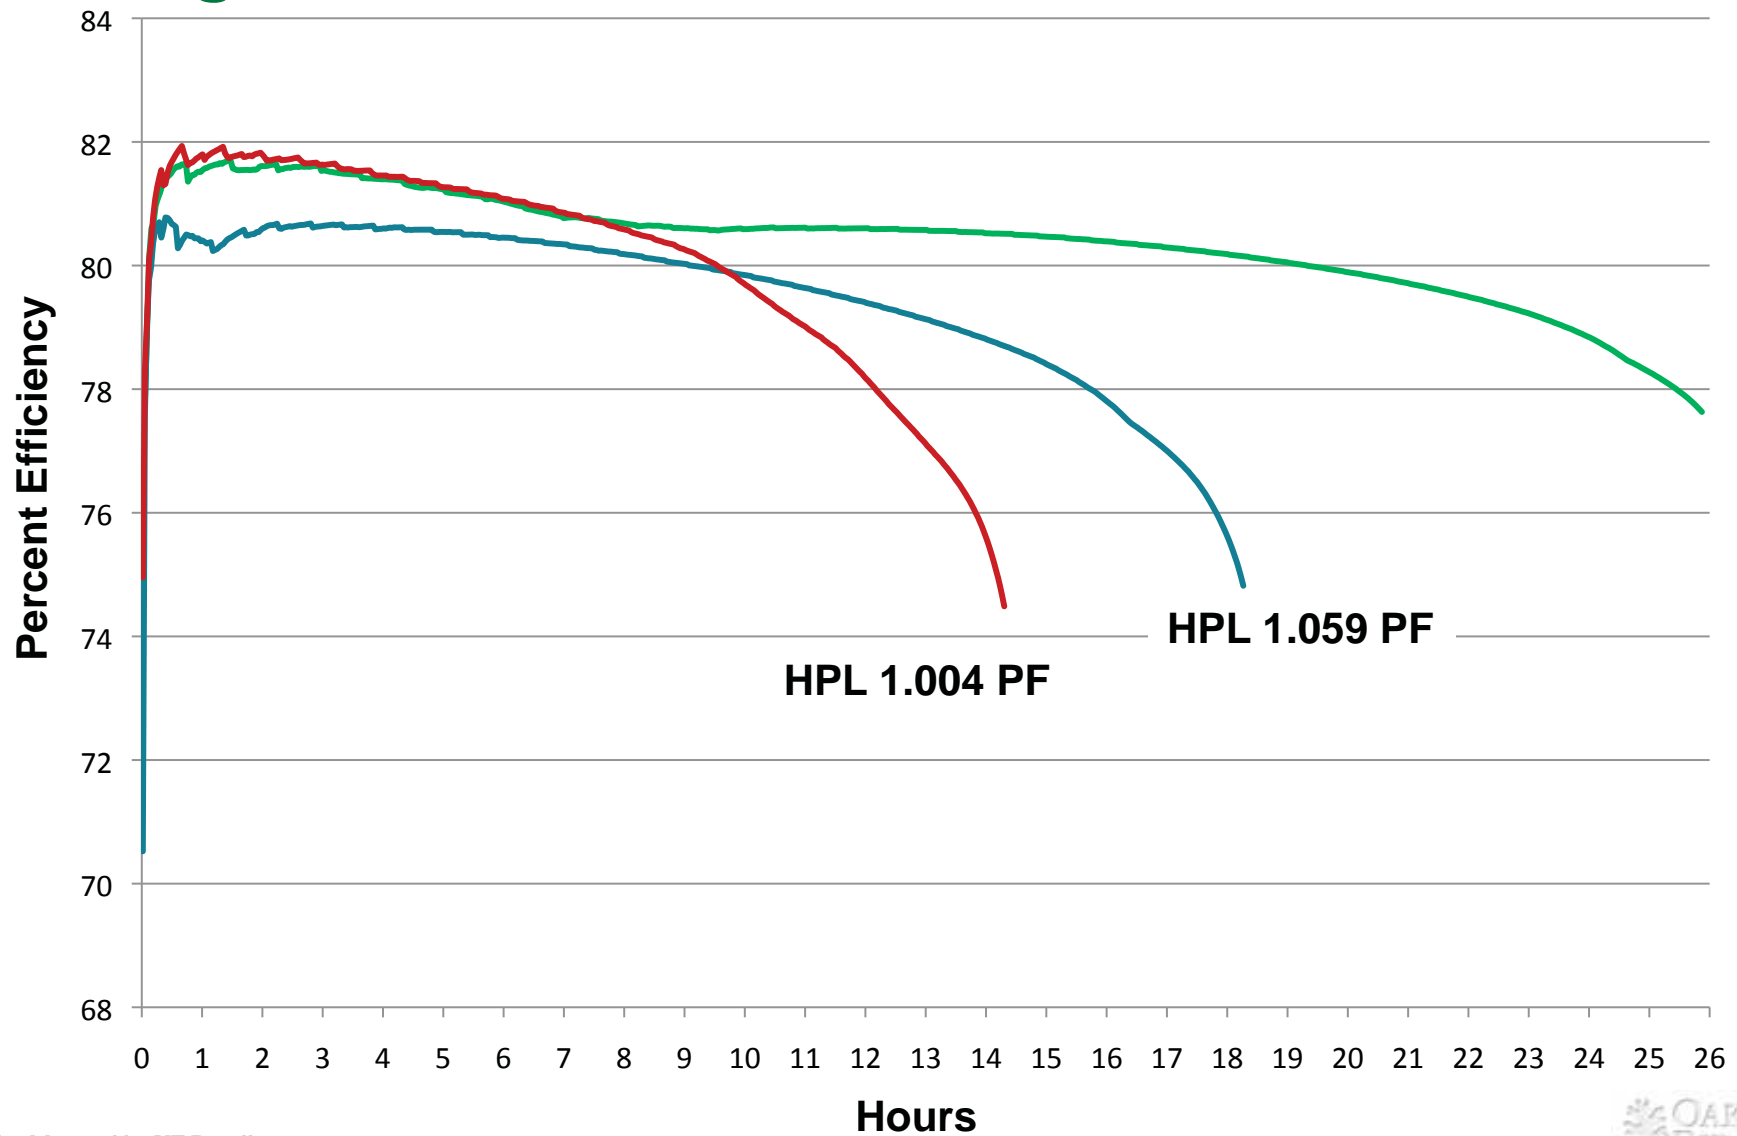
Jim Rogers jrogers@ornl.gov



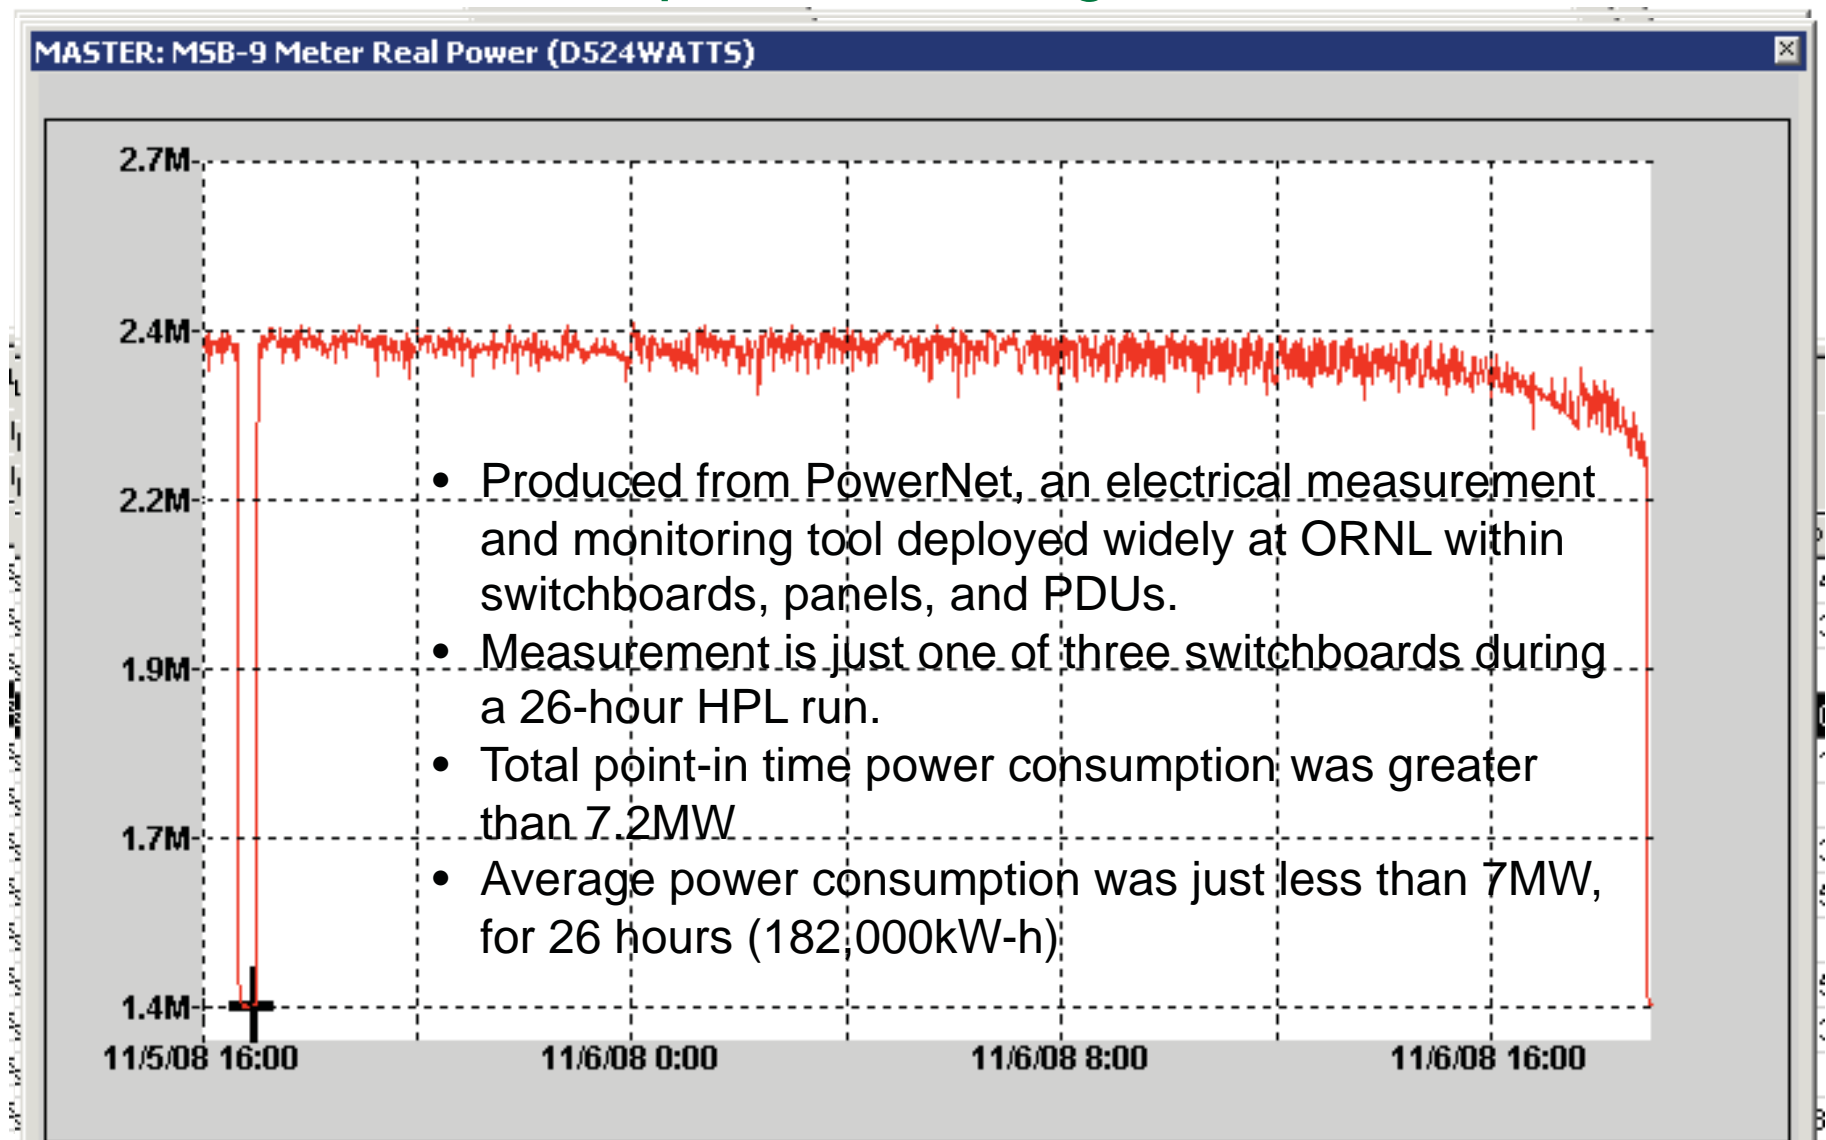
XT5 Power Allocation within a Cabinet



Jaguar- Measuring Performance Efficiency During HPL (Oct'08)




Power Consumption- Cray XT5



Cray XT5 Reliability

- 150M hours delivered through the Early Science program, and ~250M total hours delivered since the beginning of the calendar year.

- ~97% of jobs complete.
- Job failures are dominated by OOM, at more than 51%.
- Some Seastar failures are attributable to VERTY and VRM failures.
- MTTF for the Early Science period (Jan 12'09-present) is about 50 hours.



Month	Total Usage	Debug				<20 percent				>20 and <60 percent				>60 percent			
		Usage	# Jobs	Avg Job Size	% Usage	Usage	# Jobs	Avg Job Size	% Usage	Usage	# Jobs	Avg Job Size	% Usage	Usage	# Jobs	Avg Job Size	% Usage
Jan	53,602,380	0	0	0	0%	7,968,052	2,553	5,155	15%	15,613,386	482	48,786	29%	30,020,942	420	131,010	56%
Feb	52,192,949	342	1	8,000	0%	15,976,649	5,461	3,652	31%	19,333,898	477	46,606	37%	16,882,060	244	112,015	32%
Mar	69,520,641	0	0	0	0%	32,832,426	8,215	3,009	47%	26,876,482	558	42,873	39%	9,811,733	121	119,768	14%
Apr	73,124,666	1,055	4	468	0%	22,391,647	9,778	2,008	31%	33,830,940	471	42,285	46%	16,901,024	200	123,159	23%

Estimating an XT5's Operating Cost



- Curiosity- What are the anticipated operating costs (utilities) for such a system?
 - Actual kW-h measurements for operation through Mar'09 provides a good basis
 - *Assumes* that a Center needs 30% of the input power to remove the associated heat (this is a little high if you use XDP's and can measure chilled water consumption)
 - *Assumes* that a Center pays \$0.10/kW-h for electricity
 - A system the size of Kraken XT5 will have an electric bill on the order of \$2.5M.
 - A system the size of Jaguar XT5 will have an electric bill on the order of \$6-7M.

		Oct-08	Nov-08	Dec-08	Jan-09	Feb-09	Mar-09
NICS Kraken XT5	MSB8	0	27,995	810,541	1,079,763	1,011,080	1,260,445
	MSB12	0	24,930	378,807	451,061	422,821	526,089
Power Subtotal		0	52,925	1,189,348	1,530,824	1,433,901	1,786,534
Cooling Subtotal		0	15,878	356,804	459,247	430,170	535,960
Monthly Total (kW-h)		0	68,803	1,546,152	1,990,071	1,864,071	2,322,494
Cost at \$0.10/kW-h		\$ -	\$ 6,880.25	\$ 154,615.24	\$ 199,007.12	\$ 186,407.13	\$ 232,249.42
OLCF Jaguarpf XT5	MSB9	1,435,021	955,793	1,321,584	1,124,753	1,129,829	1,287,342
	MSB10	1,382,752	923,813	1,280,425	1,089,782	1,088,440	1,251,452
	MSB11	1,452,737	973,985	1,339,746	1,146,419	1,155,013	1,316,486
Power Subtotal		4,270,510	2,853,591	3,941,755	3,360,954	3,373,282	3,855,280
Cooling Subtotal		1,281,153	856,077	1,182,527	1,008,286	1,011,985	1,156,584
Monthly Total (kW-h)		5,551,663	3,709,668	5,124,282	4,369,240	4,385,267	5,011,864
Cost at \$0.10/kW-h		\$ 555,166.30	\$ 370,966.83	\$ 512,428.15	\$ 436,924.02	\$ 438,526.66	\$ 501,186.40

Noise Exposure


TABLE G-16 - PERMISSIBLE NOISE EXPOSURES (1)

Duration per day, hours	Sound level dBA slow response
8.....	90
6.....	92
4.....	95
3.....	97
2.....	100
1 1/2	102
1.....	105
1/2	110
1/4 or less.....	115

OSHA Regulation 1910.95(a)

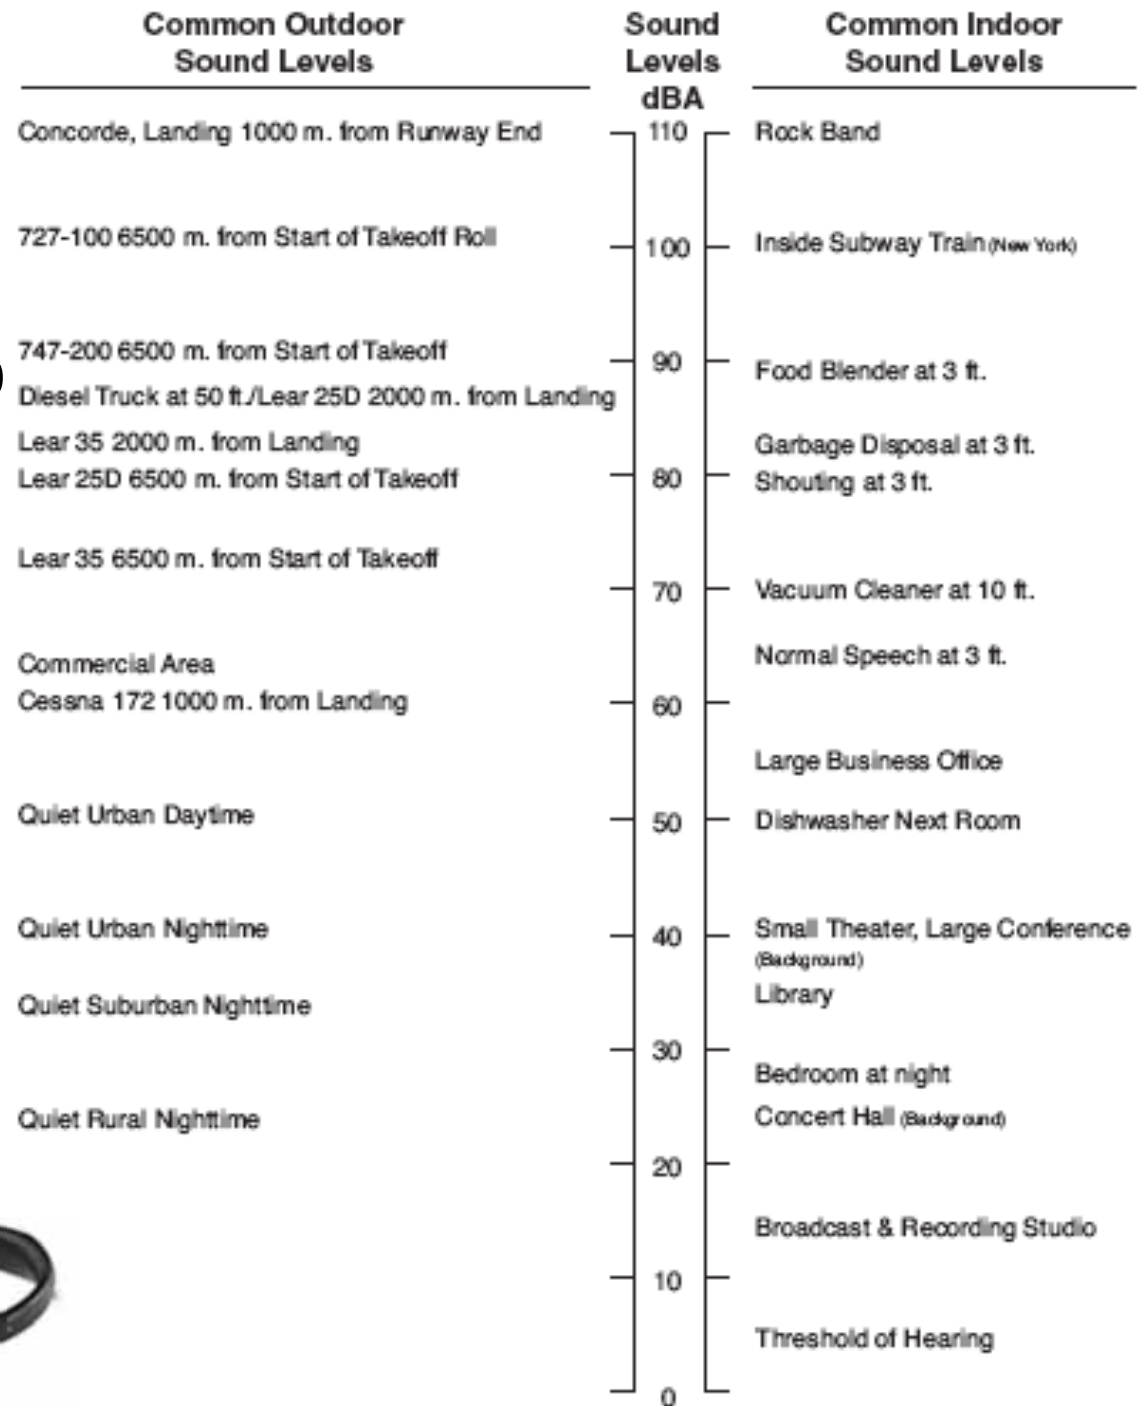
Protection against the effects of noise exposure shall be provided when the sound levels exceed those shown in Table G-16 when measured on the A scale of a standard sound level meter at slow response. When noise levels are determined by octave band analysis, the equivalent A-weighted sound level may be determined as follows:

- (translation) If the TWA (Time Weighted Average) noise level is exceeding 85 dB(A), a hearing conservation program is required
- Initial Jaguar configuration produced noise levels above 99dB
- Sound trim kits reduced this by 6-7dB
- Substantial portion of the room remains at or above 85dB, necessitating hearing protection



Noise Analogies

- XT5- Noise Operating Levels with full sound kits range from 85dB to 92dB (cabinet door open for maintenance)
- Analogous (in terms of sound pressure levels) to standing adjacent to the runway while a Boeing 747 takes off.
- Bose Quiet Comforts are not OSHA rated



INCITE April 15th call for proposals

Call for large-scale, computationally intensive, high-impact research proposals

In 2010, powerful, leadership-class computing systems at DOE's Argonne National Laboratory and Oak Ridge National Laboratory will provide over one billion processor hours to a limited number of researchers nationwide.

The call is open to scientific researchers and research organizations, including industry; DOE Sponsorship is not required. **Deadline July 1st.**

INCITE awards help advance the state-of-the-art in areas such as

- Accelerator physics
- Astrophysics
- Chemical sciences
- Climate research
- Computer science
- Engineering
- Physics
- Environmental science
- Fusion energy
- Life sciences
- Materials science
- Nuclear physics, and more

For details about the DOE leadership computing facilities, see www.alcf.anl.gov and www.nccs.gov or contact INCITE@DOEleadershipcomputing.org to be added to an announcement distribution list.

A photograph of a lecture hall with rows of metal chairs and blue seats. The chairs are arranged in a perspective that recedes into the distance. The lighting is somewhat dim, and the overall color palette is dominated by the blue of the seats and the metallic sheen of the chair frames.

Questions?