

Power Efficiency and Performance with ORNL's Cray XK7 *Titan*





Work by Jim Rogers

Presented by Arthur Bland

Director of Operations
National Center for Computational Sciences
Oak Ridge National Laboratory



ORNL's "Titan" Hybrid System: Cray XK7 with AMD Opteron and NVIDIA Tesla processors



4,352 ft²
404 m²

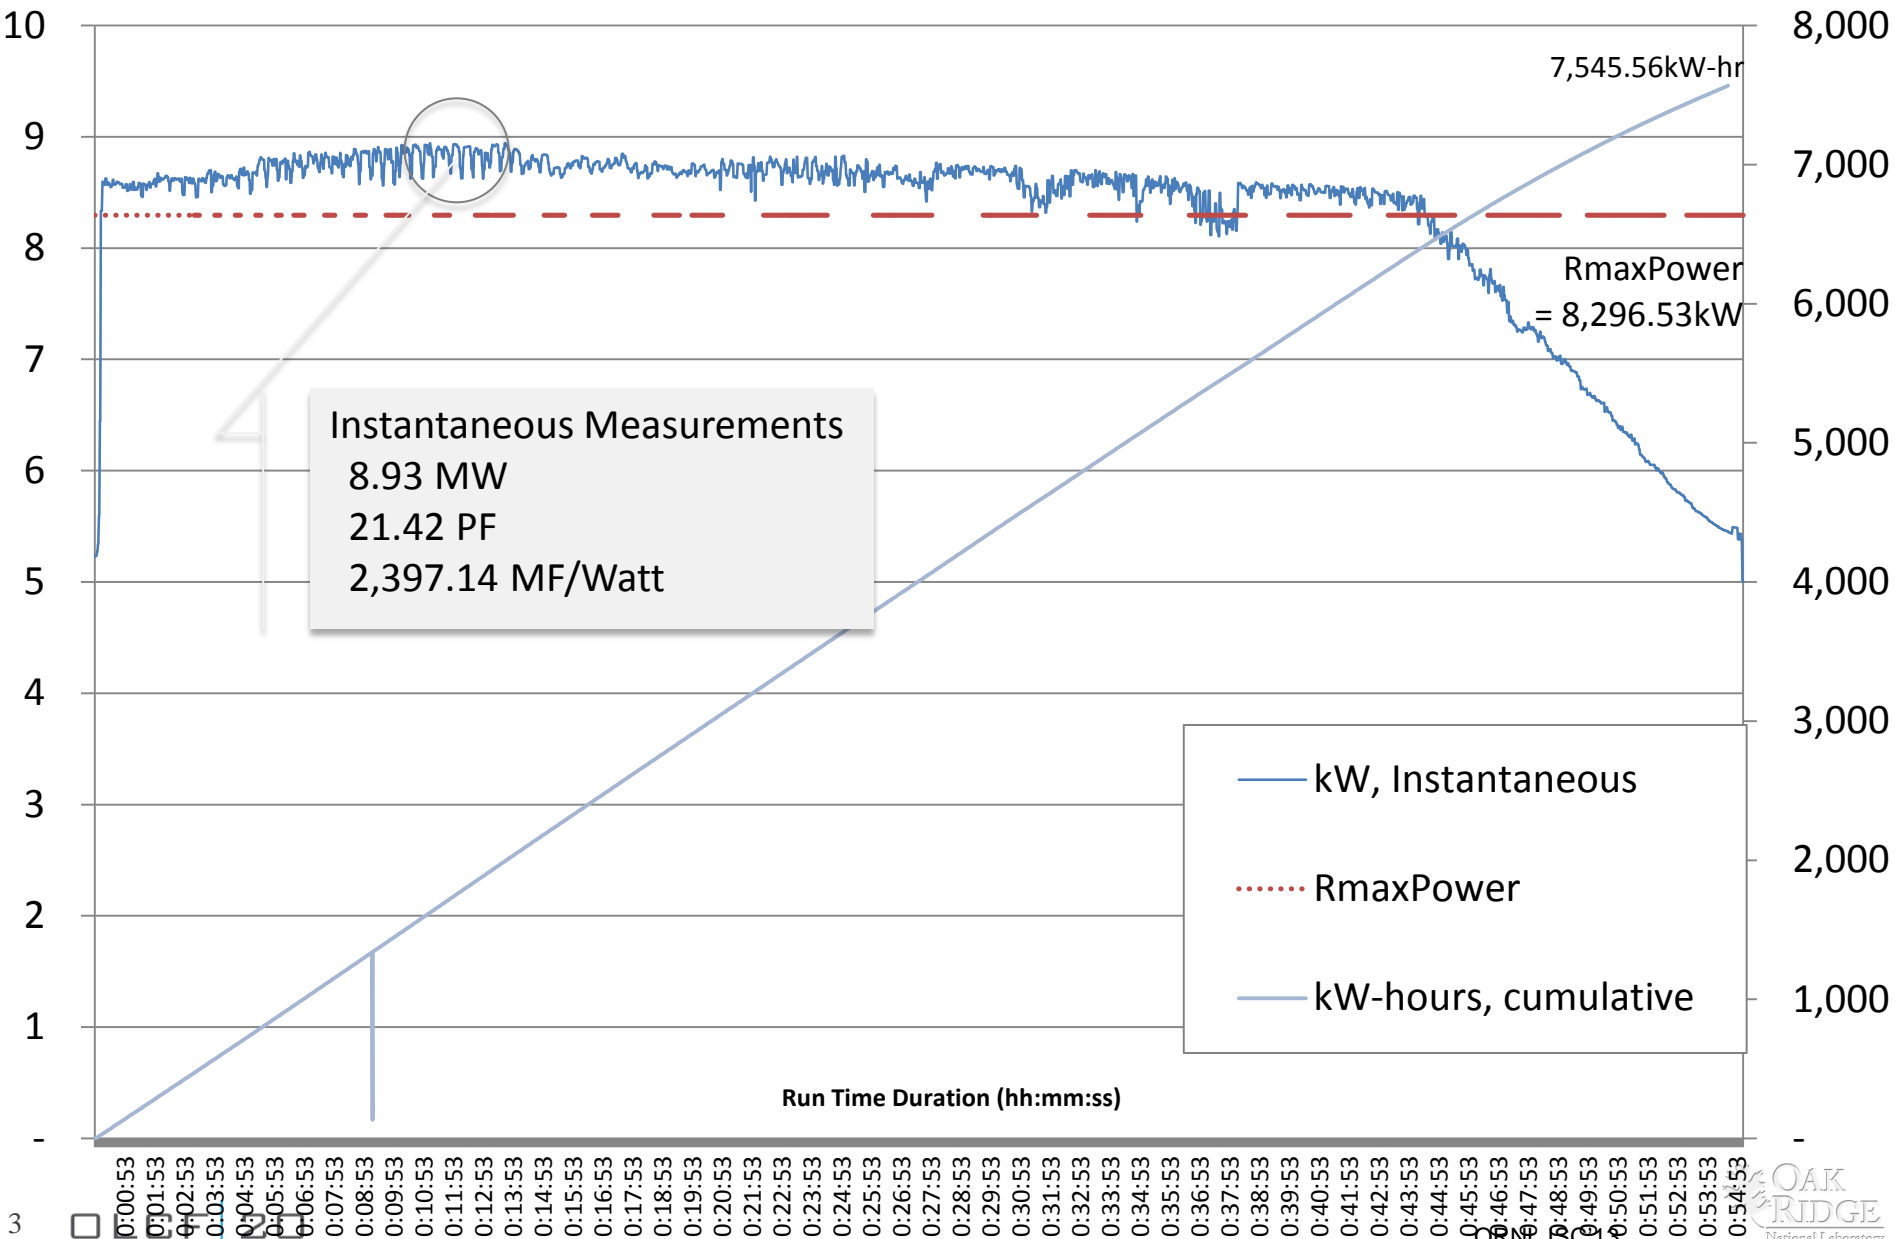
SYSTEM SPECIFICATIONS:

- Peak performance of 27.1 PF
 - 24.5 GPU + 2.6 CPU
- 18,688 Compute Nodes each with:
 - 16-Core AMD Opteron CPU
 - NVIDIA Tesla "K20x" GPU
 - 32 + 6 GB memory
- 512 Service and I/O nodes
- 200 Cabinets
- 710 TB total system memory
- Cray Gemini 3D Torus Interconnect
- 8.9 MW peak power

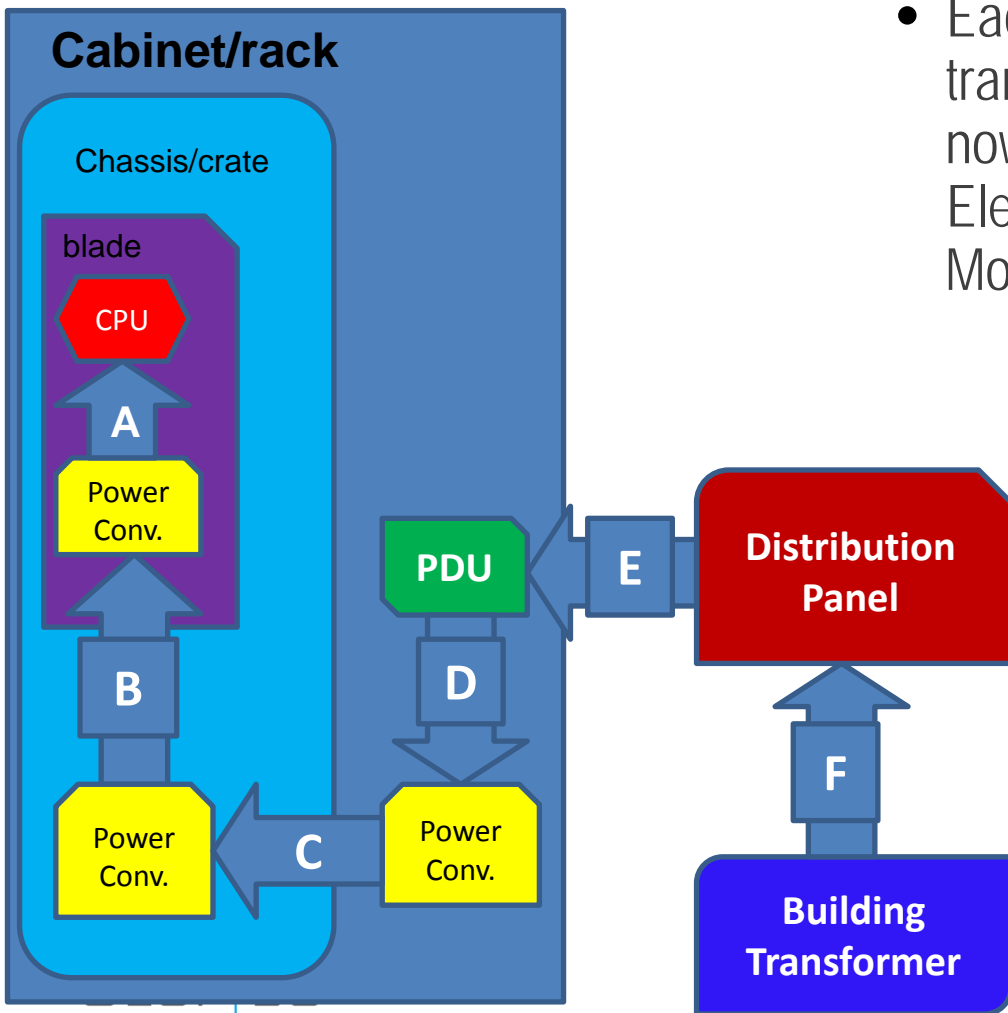
ORNL's Cray XK7 Titan | Sample Run: HPL Consumption

MW, Instantaneous

kW-hours
(Cumulative)



Revised Metering Capabilities for HPC Systems at ORNL



- Each of the seven (7) 2.5 MVA or 3.0 MVA transformers designated for HPC use is now metered by a separate Schneider Electric CM4000 PowerLogic Circuit Monitor

- "F" Position within the EEHPCWG Aspect 4 Category
- Highly accurate power quality monitor for critical energy systems. Substantially higher performance/capability than original equipment.
- Adds very accurate voltage transient and flicker analysis features.

Energy Efficient HPC System Workload Power Measurement Methodology

	Level 1	Level 2	Level 3
Aspect 1a: granularity of power measurements	1 instantaneous power sampling per second	1 instantaneous power sampling per second	continuously integrated total energy ✓
Aspect 1b: timespan of power measurements	at least one power averaged measurement covering at least 20% of the run	a time series of equally spaced power averaged measurements	a time series of equally spaced integrated total energy values ✓
Aspect 1c: reported analyzed measurements	core phase average power	core phase average power, whole application average power, idle power	core phase average power, whole application average power, idle power ✓
Aspect 2: machine fraction	the greater of 1/64 of the machine or 1 kW	the greater of 1/8 of the machine or 10 kW	whole machine ✓
Aspect 3: subsystems included	all participating subsystems, either measured or estimated	all participating subsystems, either measured or estimated	all participating subsystems must be measured ✓
Aspect 4: power measurement point	upstream of power conversion OR power conversion loss modeled with manufacturer data	upstream of power conversion OR power conversion loss modeled with off-line measurements of single power supply	upstream of power conversion OR power conversion measured simultaneously during the same run ✓

- ORNL was an early adopter/participant
- The Nov'12 Titan measurement met Level 2 Aspect 1a/1b and Level 3 aspect 2/3/4 requirements
- New meters support all Level 3 requirements for all HPC systems.