



Energy Efficiency Considerations and HPC Procurement

James H. Laros III

Principal Member of Technical Staff

Sandia Laboratories

jhlaros@sandia.gov



*Exceptional
service
in the
national
interest*



Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND2016-8123 C

Background

- Starting from scratch (2012)
 - Emerging issue for HPC
 - Most of not all efforts were still in research stage
 - Wanted to get ahead of the issue
- Goals
 1. Provide Guidance
 - Have you considered this?
 - Sites can pick and choose areas that are important to them
 2. Facilitate communication
 - Throughout HPC community
 - Between HPC community and HPC vendors
 3. Forecast community needs
 - Provide vendors some lead time to deliver solutions
- Intended to be a living document
- Introduction does a really nice job of describing what the document is and possibly more important what it is **NOT!**
 - https://eehpcwg.llnl.gov/pages/compsys_pro.htm

Approach

- Leverage existing expertise in the area
- Recognized different needs
 - System/Platform/Cabinet
 - Node
 - Component
- Express importance and forecast/predict what we will need
 - Mandatory – confident it can be delivered soon
 - Important – “think” this is reasonable in the near/mid term
 - Enhancing – what we really want even though we likely can’t get it today

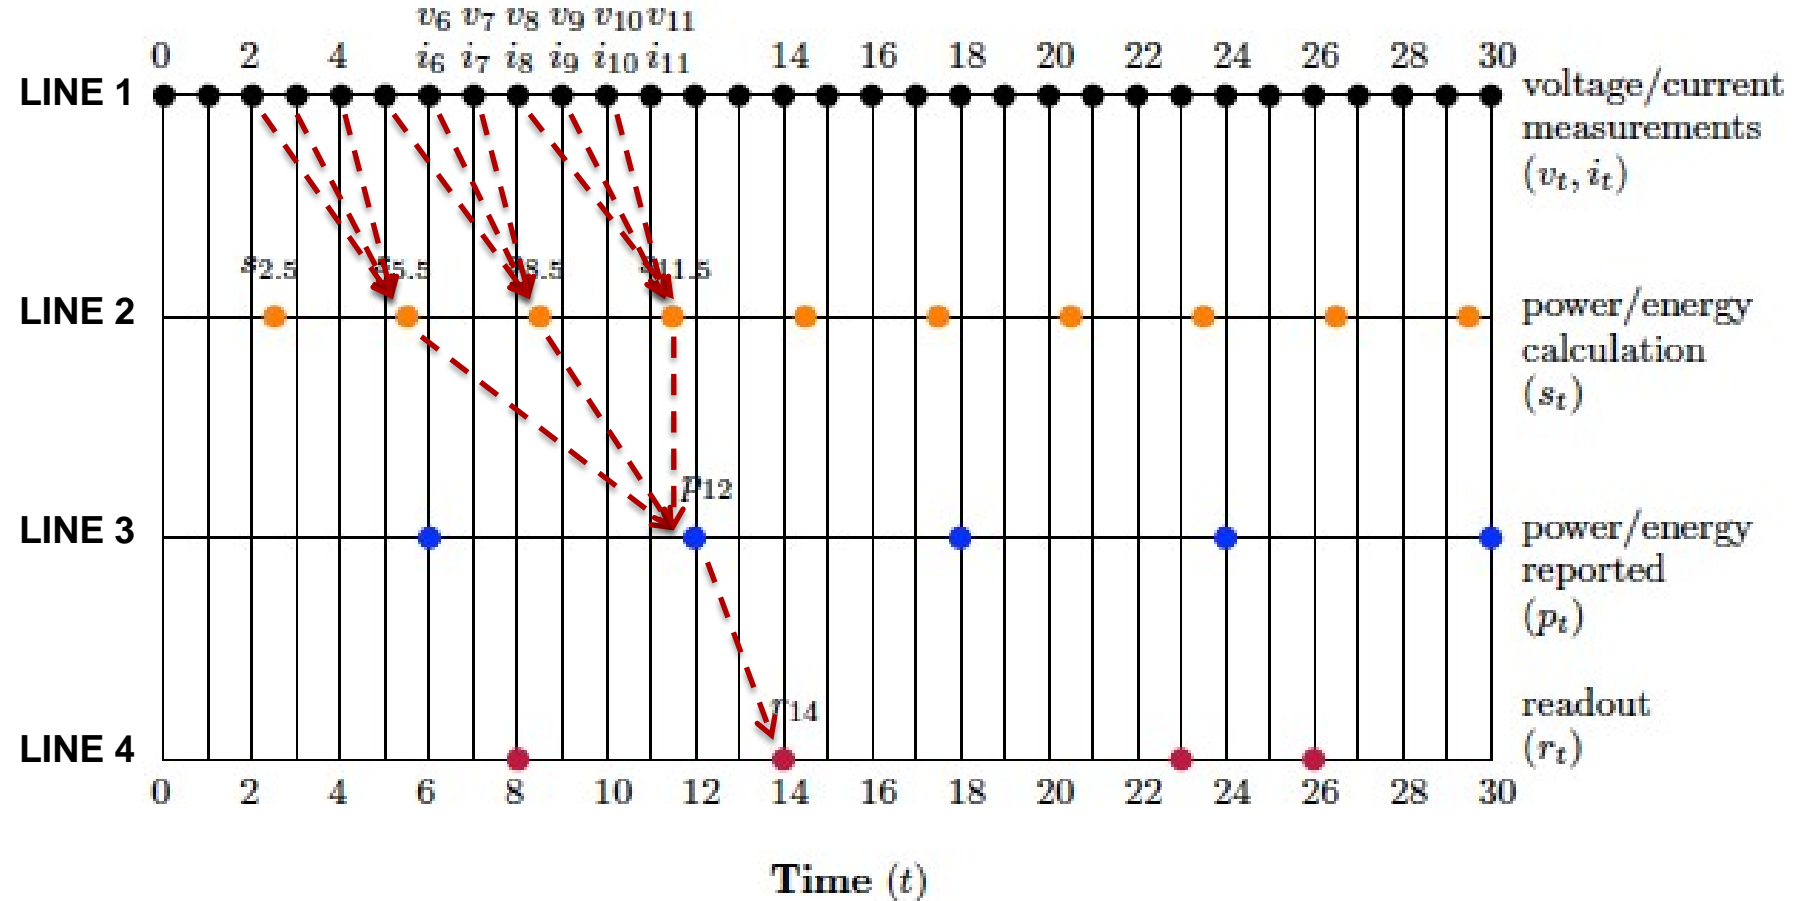
System/Platform/Cabinet

Internal Sampling Frequency	
Mandatory	≥ 10 per second
Important	≥ 100 per second
Enhancing	≥ 1000 per second

External Reported Value Frequency		
Mandatory	Discrete Power (W)	≥ 1 per second
	Average Power (W)	≥ 1 per second
	Energy (J)	≥ 1 per second
Important	Discrete Power (W)	≥ 10 per second
	Average Power (W)	≥ 1 per second
	Energy (J)	≥ 1 per second
Enhancing	Discrete Power (W)	≥ 100 per second
	Average Power (W)	≥ 1 per second
	Energy (J)	≥ 10 per second

Also Node and Component

Reported Values vs. Internal Samples



Influenced Power API Specification

SANDIA REPORT
SAND2016-4446
Unlimited Release
Printed May 2016


High Performance Computing - Power Application Programming Interface Specification Version 1.3

James H. Laros III, David DeBonis, Ryan Grant, Suzanne M. Kelly, Michael Levenhagen, Stephen Olivier, Kevin Pedretti

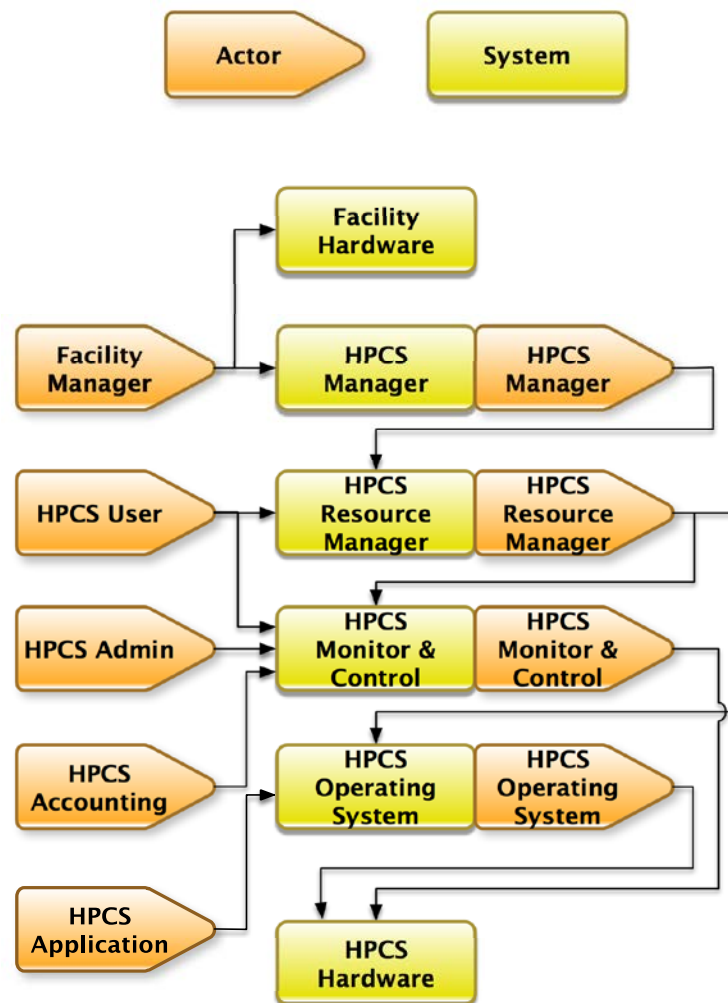
Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.



Sandia National Laboratories



Influenced Trinity RFP

Power measurement and control capabilities (hardware and software tools and application programming interfaces (APIs)) are necessary to meet the needs of future supercomputing energy and power constraints.

1. Describe all power related measurement and control features, capabilities and limitations (hardware and software) of the system including, but not limited to, any tools, system software features and APIs that will be made available at initial acceptance.
2. Describe all power related measurement and control capabilities projected on the Offeror's road map. LANS, UC, and the Subcontractor will work cooperatively to define a set of capabilities that will be delivered beyond initial acceptance.
3. Describe all power related measurement and control capabilities (hardware and software) that would necessitate hardware upgrade or replacement.

Response was used as the basis to develop an Advanced Power Management NRE program to implement the HPC Power API

Influenced Crossroads RFP

Power, energy, and temperature will be critical factors in how the APEX laboratories manage systems in this time frame and must be an **integral part** of overall Systems Operations. The solution must be well integrated into other intersecting areas (e.g., facilities, resource management, runtime systems, and applications). The APEX laboratories expect a growing number of use cases in this area that will require a vertically integrated solution.

- The Offeror shall describe all power, energy, and temperature measurement capabilities (system, rack/cabinet, board, node, component, and sub-component level) for the system, including control and response times, sampling frequency, accuracy of the data, and timestamps of the data for individual points of measurement and control.
- The system should include an integrated API for all levels of measurement and control of power relevant characteristics of the system. It is preferable that the provided API complies with the High Performance Computing Power Application Programming Interface Specification (<http://powerapi.sandia.gov>).
- And 8 more specific requirements.

https://eehpcwg.llnl.gov/pages/compsys_pro.htm

