

Exceptional service in the national interest



Power API for HPC: Standardizing Power Management and Control

Speaker: Ryan E. Grant

PowerAPI Team: James H. Laros III (Lead), Kevin Pedretti, Suzanne M. Kelly, Michael Levenhagen, David DeBonis, Stephen Olivier, Ryan E. Grant



Images courtesy of: Oak Ridge National Laboratory and Cray Inc.

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. SAND NO. 2014-15182PE

Outline


- What is the Power API?
- System Model
- Example Use Cases

What is the Power API?

- The Power API is a comprehensive system software API for interfacing with power measurement and control hardware
- Designed to be comprehensive across many different levels of a data center
- Many different actors can interface with a single API to perform several different roles
- Encompasses facility level concerns down to low level software/hardware interfaces

What is the Power API?

- Broad scoped Portable API
- Multiple actors can interact with the system at different levels
- Each interaction represents an interface that is defined in the PowerAPI



Power API System Model

The Power API System model

-Provides basic building blocks for discoverable hierarchy

Platform – only one platform can exist, encompasses whole machine

Cabinet – racks or other logical grouping of nodes, cabinet level power measurements may be supported by some hardware


Board – a piece of hardware containing multiple nodes (may not exist in all system configurations)

Node – a piece of hardware that can contain one or several sockets/cores

Socket – a physical socket that can support one or more computational cores

Core – a computational core

Other elements like memory, power planes, NICs etc can be added to the objects shown here



Common Functions

- System model allows for navigation through system objects
- Basic measurement and control via getting/setting of exposed object attributes
- Attributes for objects (node, socket, core, etc.) can be manipulated depending on role
- Metadata interface provides information about quality/frequency of exposed attributes and other information
- Statistics can be gathered on one or many object attributes over time

Who is Behind PowerAPI?

- Use case document reviewed by Labs, Universities and Commercial partners
- Use case document released as a Sandia Report in 2013
- Power API Specification reviewed by panel of experts in July 2014 at SNL
- August 2014 – Power API Specification Release (<http://powerapi.sandia.gov>)
- Meeting in Denver for public comments in Denver, September 2014



Example Use Cases

- Control power in a hardware overprovisioned system with a given MW power cap
- Accounting and prediction of power load for cooperation with power utilities
- Oversight entities wish to have long term historical power/energy data for the platform
- Users wish to monitor their jobs on fine-grained scales to understand/improve power/energy consumption
- Enables studies of whole system power/energy consumption

Thank you



<http://powerapi.sandia.gov/>



Acknowledgments:

This work was funded through the Computational Systems and Software Environment sub-program of the Advanced Simulation and Computing Program funded by the National Nuclear Security Administration