Panel on Power API/Redfish/GEOPM

Moderated by – Ram Nagappan, Intel
Panelists: Jim Laros, Sandia;
Jonathan Eastep, Intel; Ed Benson, HPE

Supercomputing Conference 2016 - Energy Efficient HPC Working Group Workshop
November 13, 2016
Salt Lake City, Utah

Panel on Power API/Redfish/GEOPM

- 10:45-11:00
  - Introduction and presentation on Power API, GEOPM, Redfish
- 11:00-11:25
  - Panel Discussion
- 11:25-11:30
  - Closure

Jim Laros

James H. Laros III is a Principal Member of Technical Staff at Sandia National Laboratories Center for Computing Research. James currently serves as the Chief Architect of Crossroads - the third Advanced Technology System deployed by the Advanced Simulation and Computing Programs National Nuclear Security Administration. James leads a number of HPC related projects at Sandia including the Advanced Architecture Test Bed project and is the lead author of the “High Performance Computing – Power Application Programming Interface Specification”.

Jonathan Eastep

Jonathan Eastep is a Principal Engineer in the Data Center Group at Intel. He earned his doctorate at MIT in Computer Science. His work at Intel spans power research as well as product development. His work and interests emphasize codesigning software, firmware, and hardware layers of a system and approaching the design process from a whole-system perspective rather than a component-centric perspective. Jonathan is the Lead Architect of the GEOPM power management framework, an upcoming open source product developed through a collaboration between Argonne and Intel. Jonathan is also Lead Architect of early power research activities for the Xeon Phi product line in the Enterprise and Government Advanced Development team.

Ed Benson

Ed Benson is a Distinguished Technologist in HPE's HPC organization. His 20+ years of HPC experience has spanned a range of technologies and roles.

He is a graduate of Tufts University

Ram Nagappan

Ram Nagappan is a System Architect in the Data Center Group at Intel. He leads power delivery and power management architecture for CORAL and next generation systems

- A **comprehensive API** for power Measurement and Control
  - **Comprehensive** = Facility to Component
  - **API** = Define the interFACE not the mechanism
- Considers all users – people and programs
  - Core (Common)
    - *Includes: Roles, Initialization, Navigation, Objects and Groups, Attributes (Get/Set), Metadata and Statistics*
  - High-Level Common
    - *Higher level of abstraction but still potentially common among multiple Roles*
  - Role/System Specific
    - Higher level abstraction specific to how Role interfaces with System

Sandia National Laboratories

http://powerapi.sandia.gov/

- Portability for the HPC community
- Forecast emerging needs of HPC community
- Expose new capabilities developed by vendors and community
  - E.g. GEOPM and Redfish
- Most important, we wanted something out there to throw stones at
- Its time for the “community” to lead this effort!
  - Version 1.0 August 2014
  - Version 1.4 October 2016
  - Version 2.0 December 2016
  - > 2.0 Community Driven

7

Global Extensible Open Power Manager

- Free open source power management framework
- Contributed to accelerate community research on power mgmt strategies to overcome Exascale challenges
- Job-level power management runtime, works in conjunction with power-aware schedulers
- Scalable, self-configuring tree-hierarchical design for Rackscale to Exascale deployments
- Plug-in architecture for extensibility in control algorithms as well as hardware platform portability
- Example plug-ins included which significantly improve performance and efficiency via application-awareness
- Will be a general product offering, working toward first deployment on CORAL Theta KNL system

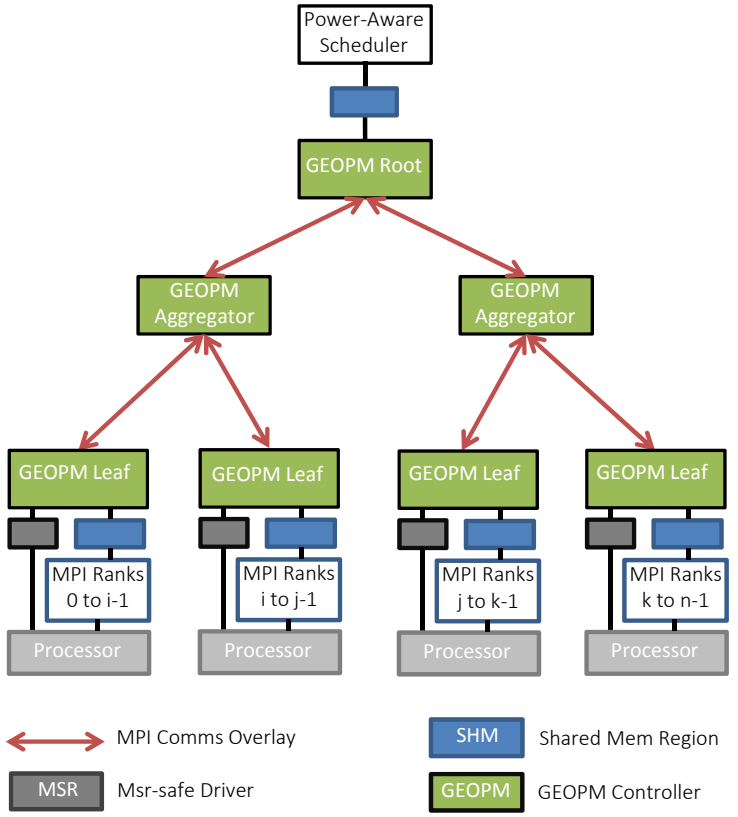

Project url: http://geopm.github.io/geopm

Synergies Between Power API, GEOPM, and Redfish

- Power API is a specification for power monitoring and control interfaces
  - Proposes common interfaces for interoperability between power mgmt implementations
- Redfish is a specification for data center management
  - Provides convenient RESTful interface for power monitoring, control, and broader data center management functions
- GEOPM is a runtime for power management
  - Implements monitoring and control, and importantly: optimizes job power/performance
  - Would sit under Power API / Redfish to implement relevant power controls and monitors
- Ongoing collaboration between GEOPM, Power API, and Redfish
  - Redfish and Power API working toward compatibility
  - Next Power API and GEOPM releases achieve near-compatibility in app-facing interfaces
- Would love to see community Power API / Redfish implementations using GEOPM

Redfish Update

- **Industry Standard specification for systems management**
  - Produced by the DMTF's Scalable Platforms Management Forum
  - RESTFul interface over HTTPS with JSON payloads
  - A secure, multi-node capable replacement for previous interfaces
- **Recent releases and activity**
  - **NEW**: Support for updating Firmware and Software (e.g. BIOS and BMC)
  - **COMING SOON**: Redfish Host Interface for OS-based access
  - **NEW**: Expanded coverage to PCIe devices, switches, fabrics, NV memory
  - **NEW**: Open Source Tools published on Github – command line interface, Redfish emulator, mockup tools, etc. www.github.com/DMTF
  - **NEW**: User Forum: www.redfishforum.com
  - **NEW**: Swordfish effort from SNIA (Released Sept 2016)
    - Extends Redfish for Storage management
- **Topics for EE HPC Working Group**
  - Seeking feedback, especially on power and thermal monitoring capabilities
  - Help with definitions of additional sensor types
  - OpenBMC effort is adding Redfish support
  - Expect a PowerAPI implementation using Redfish for all data retrieval

DMTF™

Redfish

redfish.dmtf.org

GitHub

SNIA

Swordfish™

www.snia.org/swordfish

Power API with Redfish Prototype Example

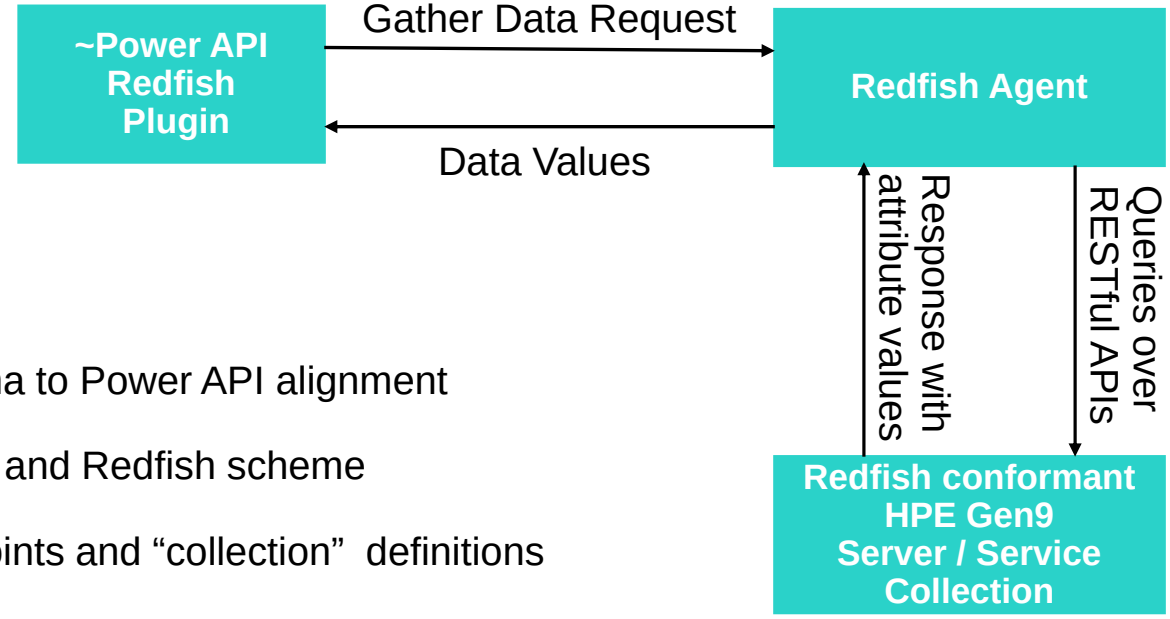

Examining:

Redfish Schema to Power API alignment

HPE's sensors and Redfish scheme

Aggregation points and "collection" definitions

Scaling…….

Panel Discussion & Questions

Lunch Discussion Group

13:00-13:45 - Power API, GEOPM and Redfish Discussion Group

Chat informally with panelists and moderator

Backup

Panel Discussion/Questions

- Current deployment of Power API, GEOPM and Redfish
  - What went well?
  - What are the challenges?
- Can we reduce operational expenditures and improve energy efficiency using these technologies?
- How do you see these technologies help pre Exascale and Exascale systems?
- Is this only applicable for very large Systems (Ex: Top 10)?
  - How this will help smaller systems?
- Is this open source? How do I get this?