

# Energy Management IBM

Todd Rosedahl, Chief Engineer

IBM/POWER Firmware Development

# IBM Customer Sets

- **Big Business - Banks, Wall Street, Insurance**
  - Traditionally focused on Performance only
    - Some starting to see the advantage of active off-peak management
    - TCO, which includes power consumption, sometimes used at buy-time
    - Some customers do need accurate energy estimates pre-purchase
  - Density, trapped/stranded power, and cost of energy not a concern
  - Focus on making money from their core business
  - Virtualization/Consolidation has been the main play
- **Small Business/Satellite offices**
  - Concerned about noise, but not energy consumption
- **Governments (US and Foreign)**
  - EnergyStar (or similar) standard must be met
  - Idle power is most crucial
  - SERT (active workload) emerging
- **DataCenters/Cloud Providers - TenCent, Rackspace, Google, Softlayer**
  - Core business is providing computation
  - Density (nodes/rack) critical
    - Effective Power Capping required to reduced trapped capacity at rack level
    - Statistical Risks are taken
  - Common Interfaces Required
  - In some cases (known workload), Perf/Watt is critical
- **HPC – IBM POWER has very little of the HPC market**
  - Power consumption, ramp up/down, consistent node to node performance

# Use Case Example

Instructions: Describe a use case in a production environment for how one of your customers has leveraged your product's power and energy measurement and management features. What did they do? What was the impact?

- Customer: Large cloud provider
- Problem: Maximize node density and minimize trapped power
  - 6KW Racks that must hold 8 systems per the bid
  - As configured, we estimated a max of 900W/system
  - Limited density to 6 systems/rack (5400W)
  - Actual power consumption was normally less than 750W with common workloads
  - 1500W unused, Required density was not met.
    - 6 nodes at 750W = 4500W. 6000W – 4500W = 1500W.
- Solution: Power cap the systems
  - We capped each system to 750W
    - Did require a firmware change to enable such a low cap for this special configuration
  - All 8 systems fit and 6KW used effectively

# Q/A

- What are the dominant use cases that you find in the production environment today?
  - Performance Boost and Power Capping are the two main uses.
- What layer of the software stack do you think will mature the most quickly for power and energy management within the next two years? Why?
  - BMC
  - Low level firmware (OCC, EC, embedded device firmware)
  - Hypervisors
  - Partition firmware
  - o/s
  - Applications
  - **Cluster/Cloud/systems management**
    - **As everything goes cloud, software that dispatches work efficiently to the best node and powers nodes/cores/resources on/off as needed will become a large factor**
- What kind of energy efficiency metrics have you seen used in RFP documents?
  - Total power consumption is sometimes in the bids as well as power capped values
  - Energystar is in all government bids
  - Never seen SERT or SPECPOWER or any PPW in any bid
- Trapped and stranded capacities are concepts that suggest particular use cases. Are your customers asking for capabilities to address this? How does your product help centers to manage trapped and stranded capacity effectively?
  - Energy Estimators for stranded power
  - Power capping for trapped power

# Wish List

- **Standardized interfaces at every level**
  - Specification must have broad and strict compliance requirements
  - BMC, o/s, PDU, Power supplies, VRMs, fans, VPD, temp sensors(PCI), etc
- **Knowledge of customer power/performance workload data**
  - Eg. Memory bandwidth vs memory power consumption
  - Integrated and standardized performance/power monitoring tooling?
    - Must be vendor accessible
- **Useful (and used) PPW benchmarks. Power included in every benchmark.**
  - Memory intensive, CPU intensive, IO, etc
- **Standardized active and idle power controls on every component**
  - GPUs, FPGAs, SSD, PCI