Information Technology Industry Council

Innovation. Insight. Influence.

Server Technology Considerations
ENERGY STAR®

June 25, 2014

Henry ML Wong,
Sr. Staff Technologist
Henry.L.Wong@intel.com, 408-765-5047

Future Technology Considerations

Background

Server features and technologies continue to evolve to maximize effective output within an energy envelope.

Data demands, compute requests, availability and reliability continue to grow.

Continued technology challenge provide greater and greater output while mitigating energy and demands.

Nota Bene:

- The technologies indicated may vary in maturity and adoption in the industry.
- The technologies are presented for discussion purposes to consider as the ENERGY STAR team investigates v3.0
- As a forward looking review, actual adoption or the implied benefits may not materialize either in part or in its entirety
- New technologies derived from similar concepts may arise to nullify the adoption or integration of the technologies mentioned.
- Though various product features may have implemented something similar, the enclosed descriptions are generic concepts. As such they may not be a full representation of the concept as other solutions may appear.

Upcoming Server Technologies That May Change Server Power Profiles (cont.)

Local data storage (in the server or compute enclosures)

- SSD: solid state drives and hybrid variants
- Interfaces: SATA 3, SCSI, PCIe
- Local data hierarchy (e.g. cache levels, DRAM, Solid state, rotational, alternate technologies)
- Compression and Security

I/O (cards transition to "on board")

- System to system "Fabric"
- Wired 10Gbe build out
- Wired 40Gbe+ (ganged ports)
- Optical

Upcoming Server Technologies That May Change Server Power Profiles (cont.)

Memory (aka main system memory)

- DRAM: speed/technology, channels/buffers
- Non-volitile/solid state: affordability, and speed
- Impacts due to hierachy
- RAS/security- logical partitioning impacting dynamic load

CPU (or compute structure)

- Market application optimization, e.g. microserver, HPC
- Customization, e.g FPGA, Heterogeneous cores, etc..

Auxiliary Processing Accelerators

Auxiliary Processing Accelerators

- Additional processing units either via cards, mounted on-board, or integrated to the processor to accelerate compute functions.
- Typically complete logical systems (e.g. memory/compute) designed for large intensive processing requirements.
- "Big Data", "IoT", and "Smart xxx" in addition to traditional High Performance Computing applications has created a greater demand for APA's compute capabilities.
- As part of customization to workloads, many industry stakeholders believe the demand for such capability will increase lending to more onboard/integrated solutions
- As with HPC workloads, the utilization and compute intensity is very high. It's unlikely typical power management methods have meaningful effect on these subsystems.

* IoT: Internet of Things- term used to described the vast amount of components and systems that are being linked and integrated to the internet
* Big Data: large unstructured data being collected, synthesized and analyzed. Generally used in reference to the internet.
* Smart xxx – to describe an integrated and intelligent network of systems to enable efficient infrastructure, e.g. Smart Cities

Thermal Reporting & Considerations

Thermal Reporting and Considerations

Discussion points:

- Exhaust temperatures for heavily loaded systems are 54-60 C irrespective of inlet temperature.
- Not always the case that max exhaust temp is under worst case conditions (e.g. fan speed controls)
- Test conditions, (e.g. Inlet temperature, workload) may require more precise definition to apply to data center thermal conditions

Thermal Information *	Low End	Minimum	Typical	Maximum	High End
Total Power Dissipation (watts)					
Delta Temperature at Exhaust at Peak Temp. (°C)					
Airflow at Maximum Fan Speed (CFM) at Peak Temp.					
Airflow at Nominal Fan Speed (CFM) at Nominal Temp.					

* References: ASHRAE Extended Environmental Envelope Final August 1, 2008
Thermal Guidelines for Data Processing Environments, ASHRAE, 2004, ISBN 1-931862-43-5
Peak temperature is defined as 35 °C, Nominal Temperature is defined as 18 - 27 °C

Status

- ASHRAE reviewing ENERGY STAR thermal data requirements in upcoming meeting in Seattle, WA.
- Proposal to convene an ad hoc committee from ASHRAE TC 9.9 will be put forth with representatives from manufacturers, data center owner/operators and design firms at the Seattle meeting.
- The expectation is that a supportable proposal for IT manufacturers will be created that comprehends reporting of environmental data (e.g. inlet temperature, airflow, power consumption, etc.) for use/integration with data center power/thermal designs.

Cold Redundancy

Cold Redundancy

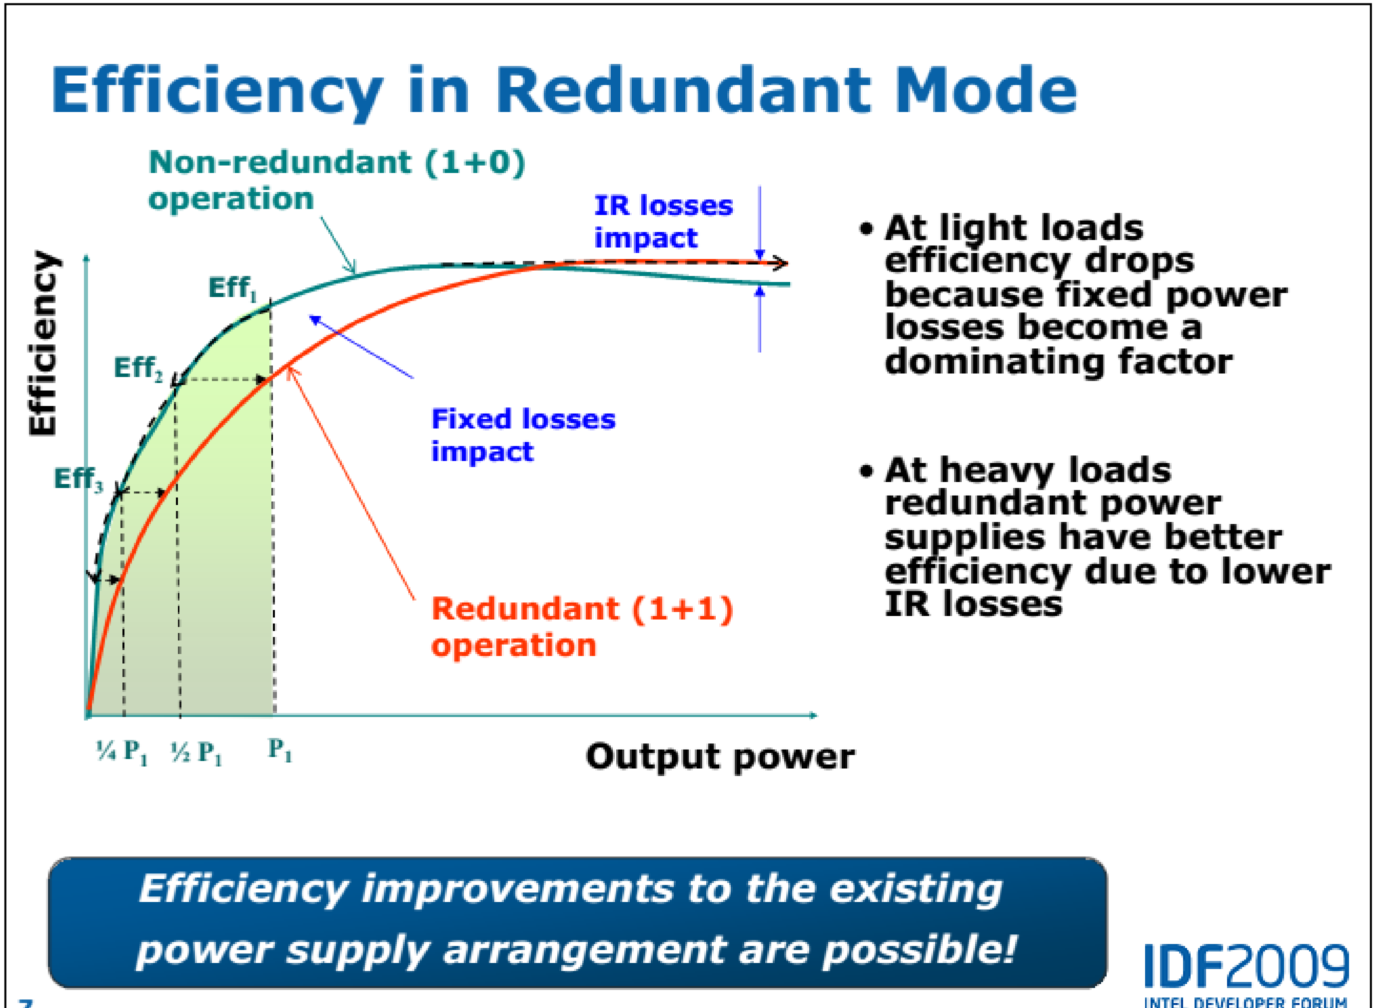

Courtesy of Intel® Corp, V. Vogman and A. Watts, Intel Developer's Forum 2009

Cold Redundancy

Cold Redundancy Technology

- Based on system power levels, the control circuitry activates only the number of power supply modules needed in order to provide maximum efficiency
- 1+1 configuration: only one power supply is active and the redundant power supply is put into a standby state
- The control circuitry continuously monitors the condition of the active power supply and enables the redundant power supply when:
  - The system goes into a high power state
  - A fault occurs in the active power supply

Courtesy of Intel® Corp, V. Vogman and A. Watts, Intel Developer's Forum 2009

Cold Redundancy

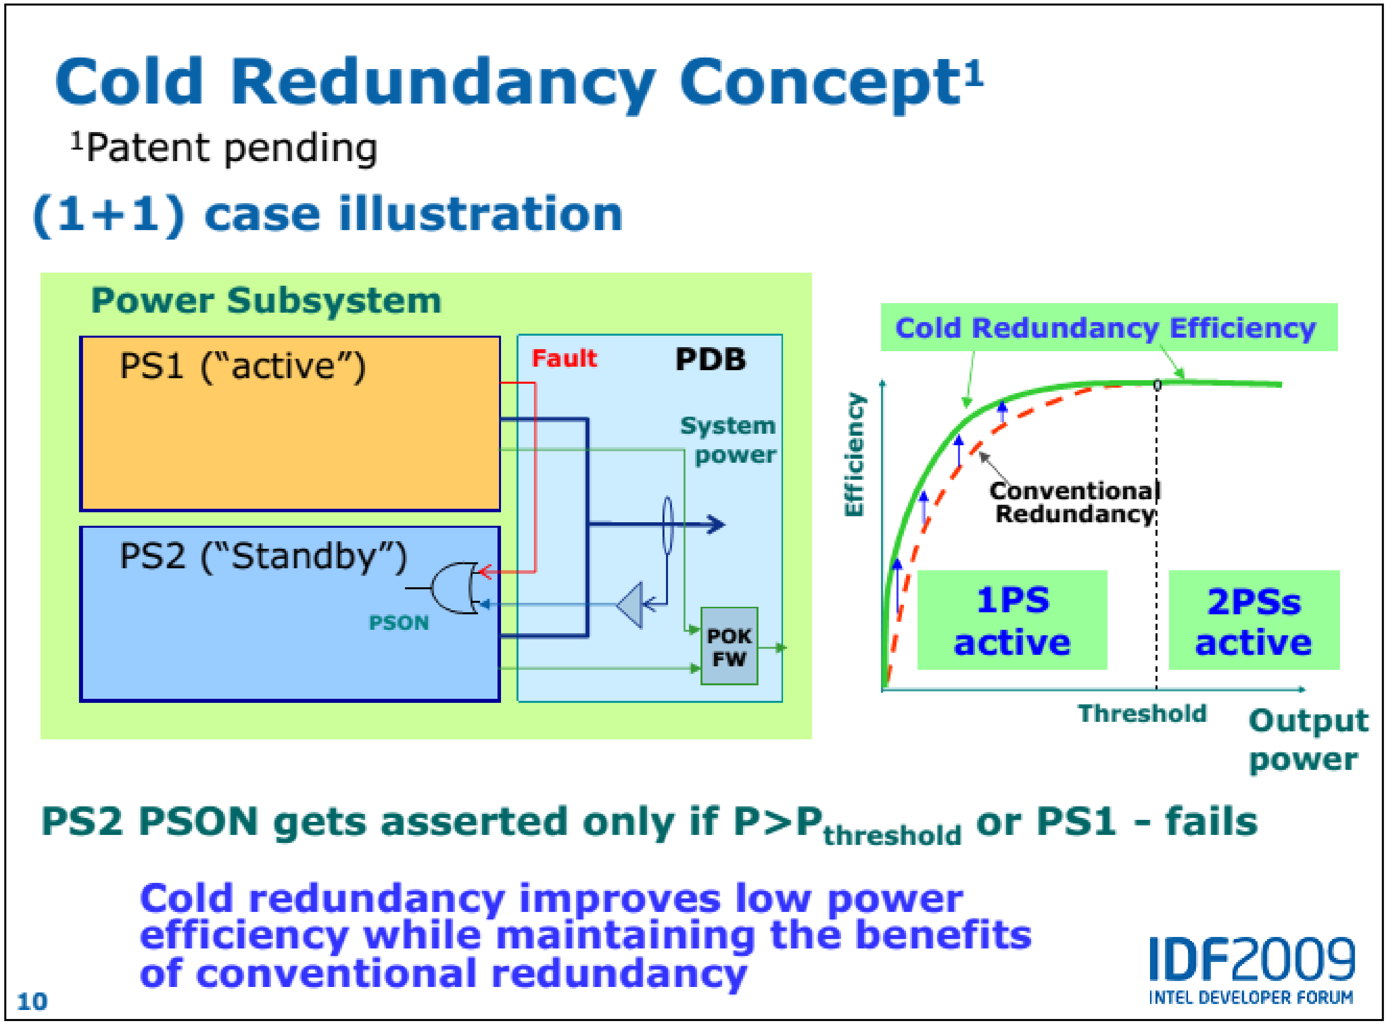

Courtesy of Intel® Corp, V. Vogman and A. Watts, Intel Developer's Forum 2009

Actual Improvements (Idle Power Range)

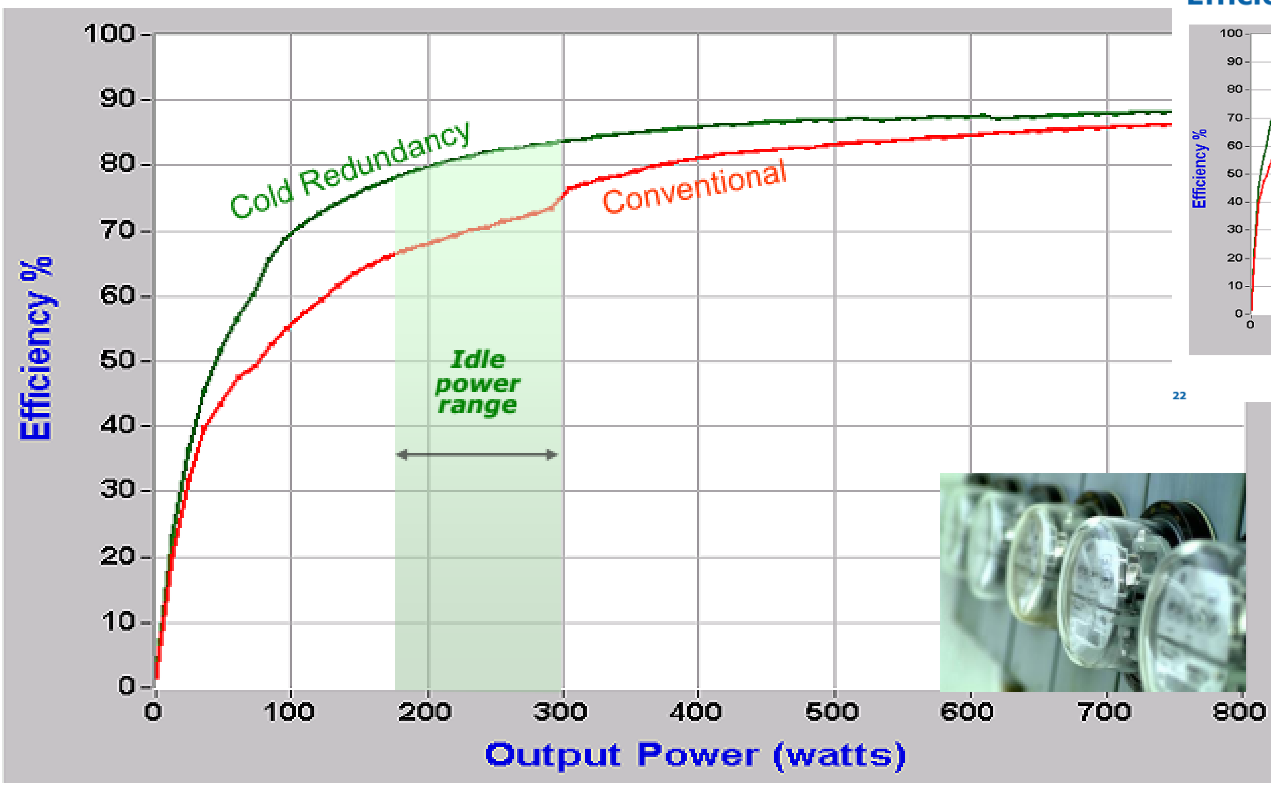

Efficiency Graphs

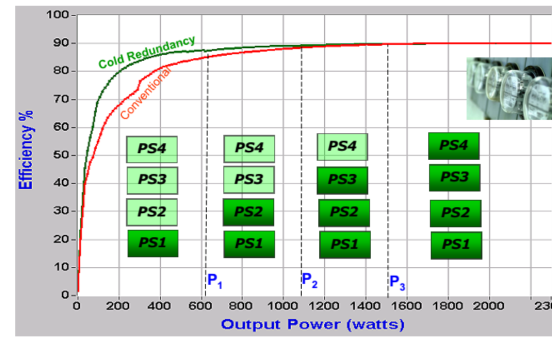

IDF2009 INTEL DEVELOPER FORUM

22

IDF2009 INTEL DEVELOPER FORUM

Courtesy of Intel® Corp, V. Vogman and A. Watts, Intel Developer's Forum 2009