AMD Opteron APU Energy Efficiency

ENERGY STAR Computer Servers Off-Season Meeting

June 2014

David Reiner | Server Performance & Power Optimization Manager

Donna Sadowy | Sr. Manager, Government Relations

AMD 2013-2014 SERVER ROADMAP

AMD

	2013	2014
2P and 4P Enterprise, Mainstream Platforms	**AMD Opteron™ 6300 and 4300 Series** 4, 6, 8, 12 or 16 “Piledriver” CPU Cores 35W-140W 32nm	**“Warsaw” CPU** 12 or 16 “Piledriver” CPU Cores 32nm
1P Web/Enterprise Services Clusters	**AMD Opteron™ 3300 Series** 4 or 8 “Piledriver” CPU Cores 25W-65W TDP 32nm	**“Berlin” CPU/APU** 4 “Steamroller” CPU Cores GCN Graphics Compute Units (APU) HSA Features (APU) 28nm
	AMD Opteron™ X1150 CPU and X2150 APU 4 “Jaguar” CPU Cores GCN Graphics Compute Units (APU) 9W-22W 28nm SoC	**“Seattle” CPU** ARM “A57” CPU Cores 28nm SoC

AMD roadmaps are subject to change without notice or obligations to notify of changes. Placement of boxes intended to represent first year of production shipments.

APU PROVIDES GREAT ENERGY EFFICIENCY

INTEGRATION

- Combines CPUs and GPU onto a single die
- APU energy performance optimization
  - Efficient, fine-grained power management between CPU and GPU
  - CPU<->GPU communication power dramatically reduced relative to separate chips
  - Shared memory interface helps save power

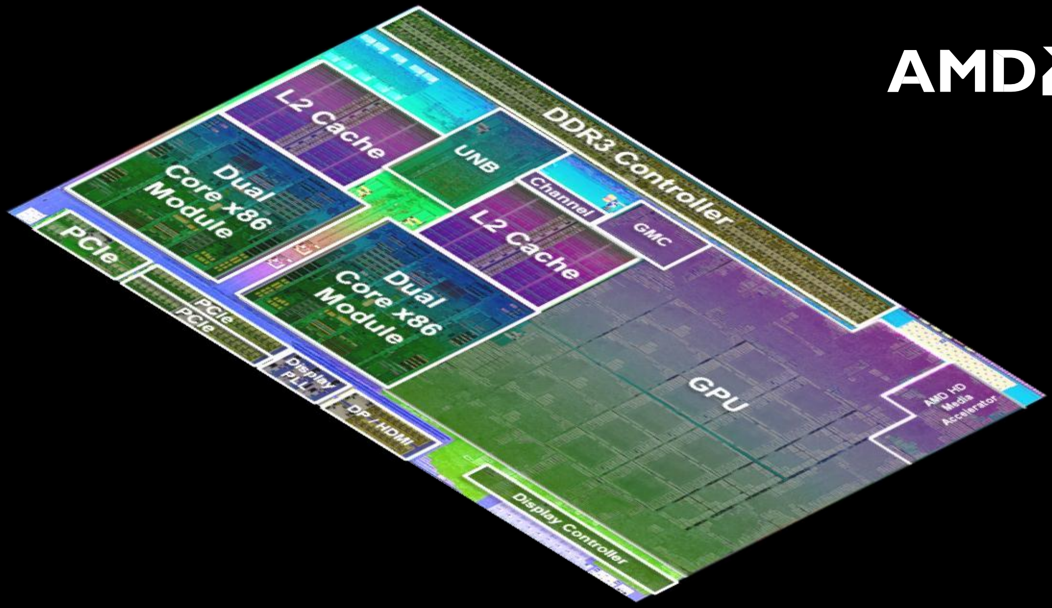

Temperature distribution plots

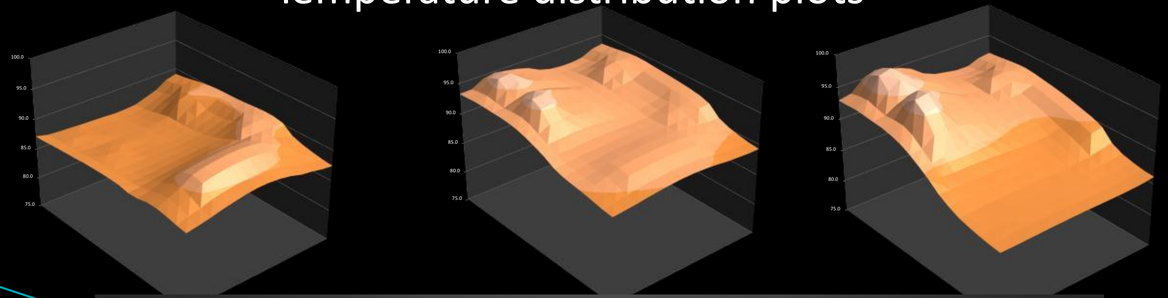

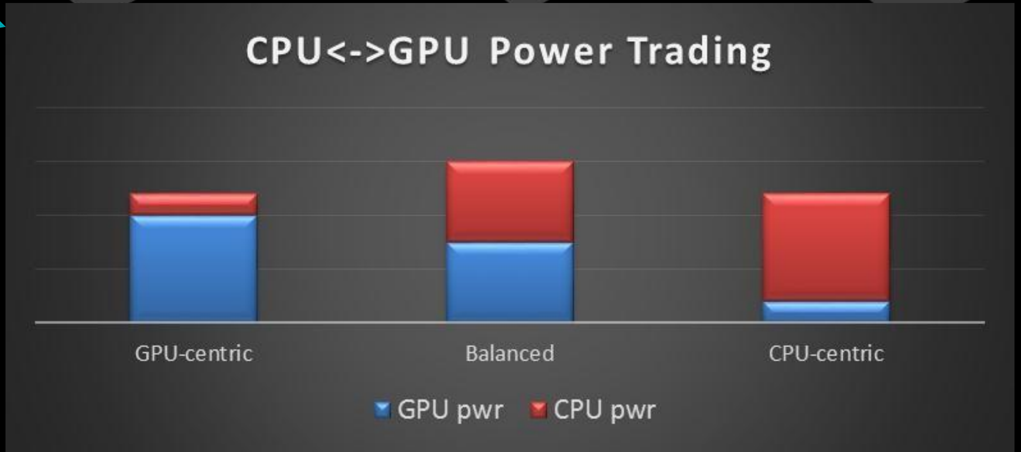

NOT ONLY HARDWARE, BUT SOFTWARE TOO

AMD

- HSA is a new standard that allows different compute elements to be mixed together on the same piece of silicon
  - Example: CPU, GPU and DSP
- Energy efficient – running parallel code on the GPU
  - CPU and GPU share main memory eliminating processor cycles to move between the two
  - Reduced instructions path also minimizes cycles
- HSA enables work to be routed to best resource
  - Some cores optimized for:
    - High Throughput
    - Low latency
    - Low Power
    - Special Accelerators
- Net effect: reduces cycles and power consumption

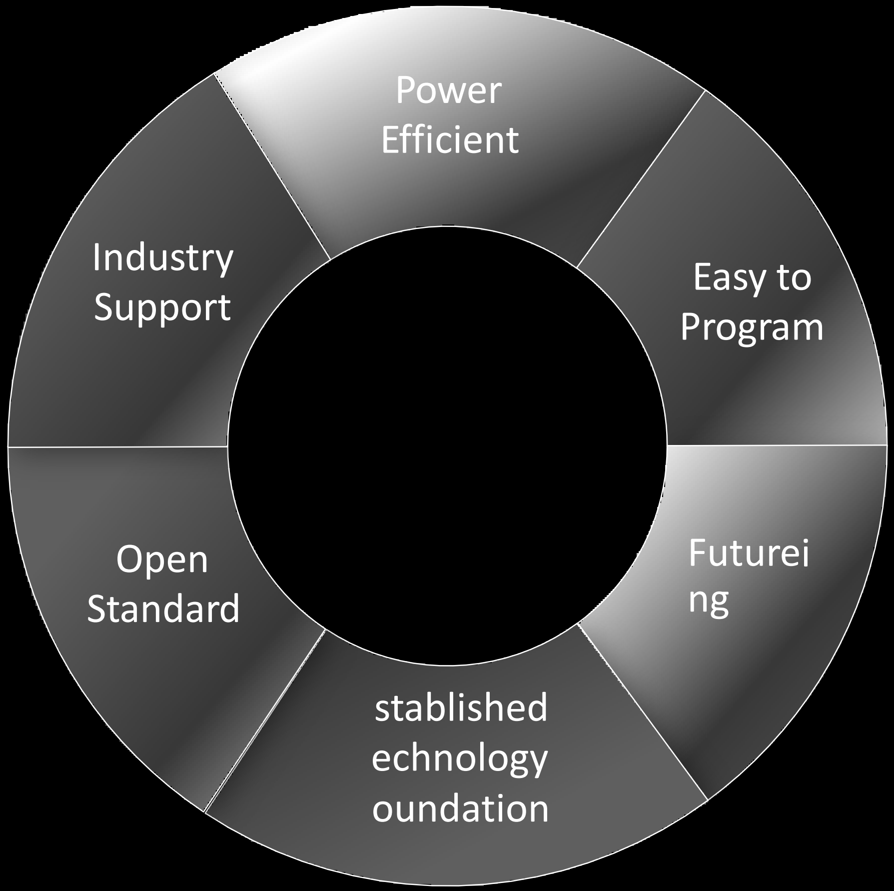

What is HSA?

Joins CPUs, GPUs, and accelerators into a unified computing framework

Single address space accessible to avoid the overhead of data copying

Use-space queuing to minimize communication overhead

Pre-emptive context switching for better quality of service

Simplified programming

Single, standard computing environments

Support for mainstream languages—C, C++, Fortran, Java, .NET

Lower development costs

Optimized compute density

Radical performance improvement for HPC, big data, and multimedia workloads

Low power to maximize performance per watt

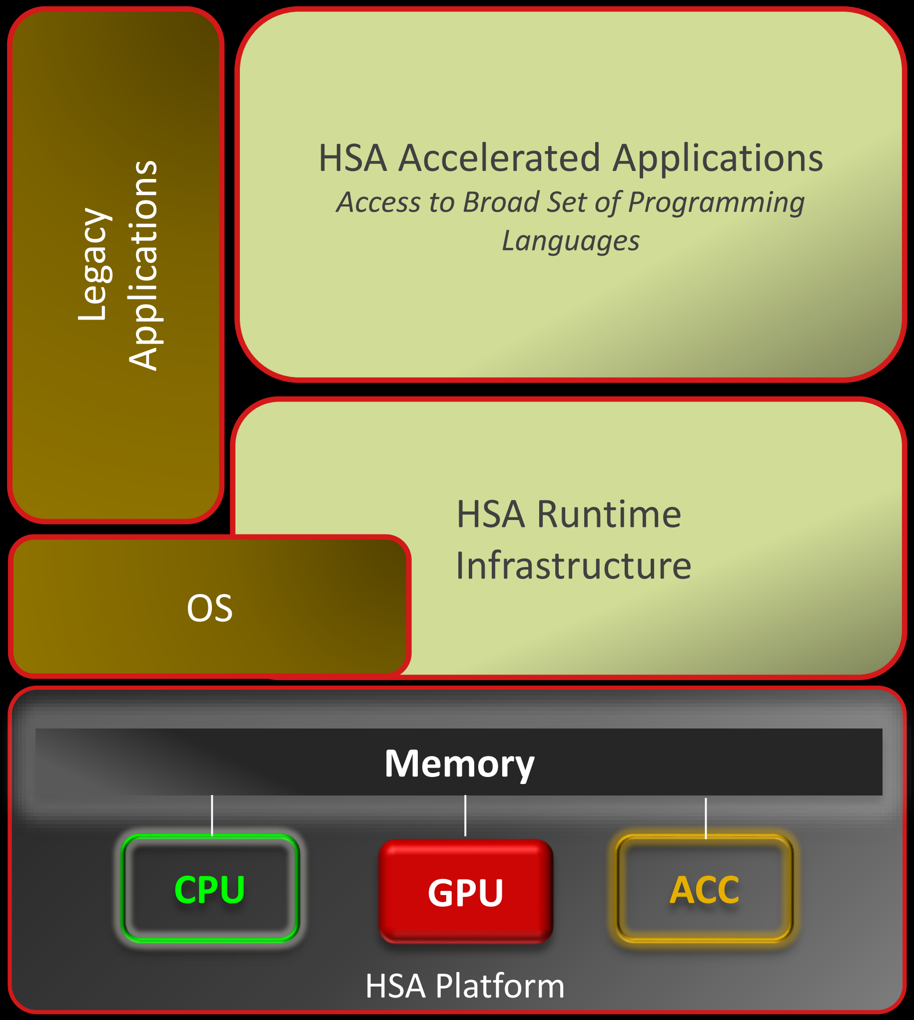

FUTURE ENERGY EFFICIENCY GAINS

AMD

- Continued innovation in power management, integration, IP efficiency will reduce average energy use
- Accelerated CPU performance gains coupled with HSA enabled GPU compute will drive overall APU performance
- Fewer battery charges, less infrastructure material (i.e. thinner and lighter)

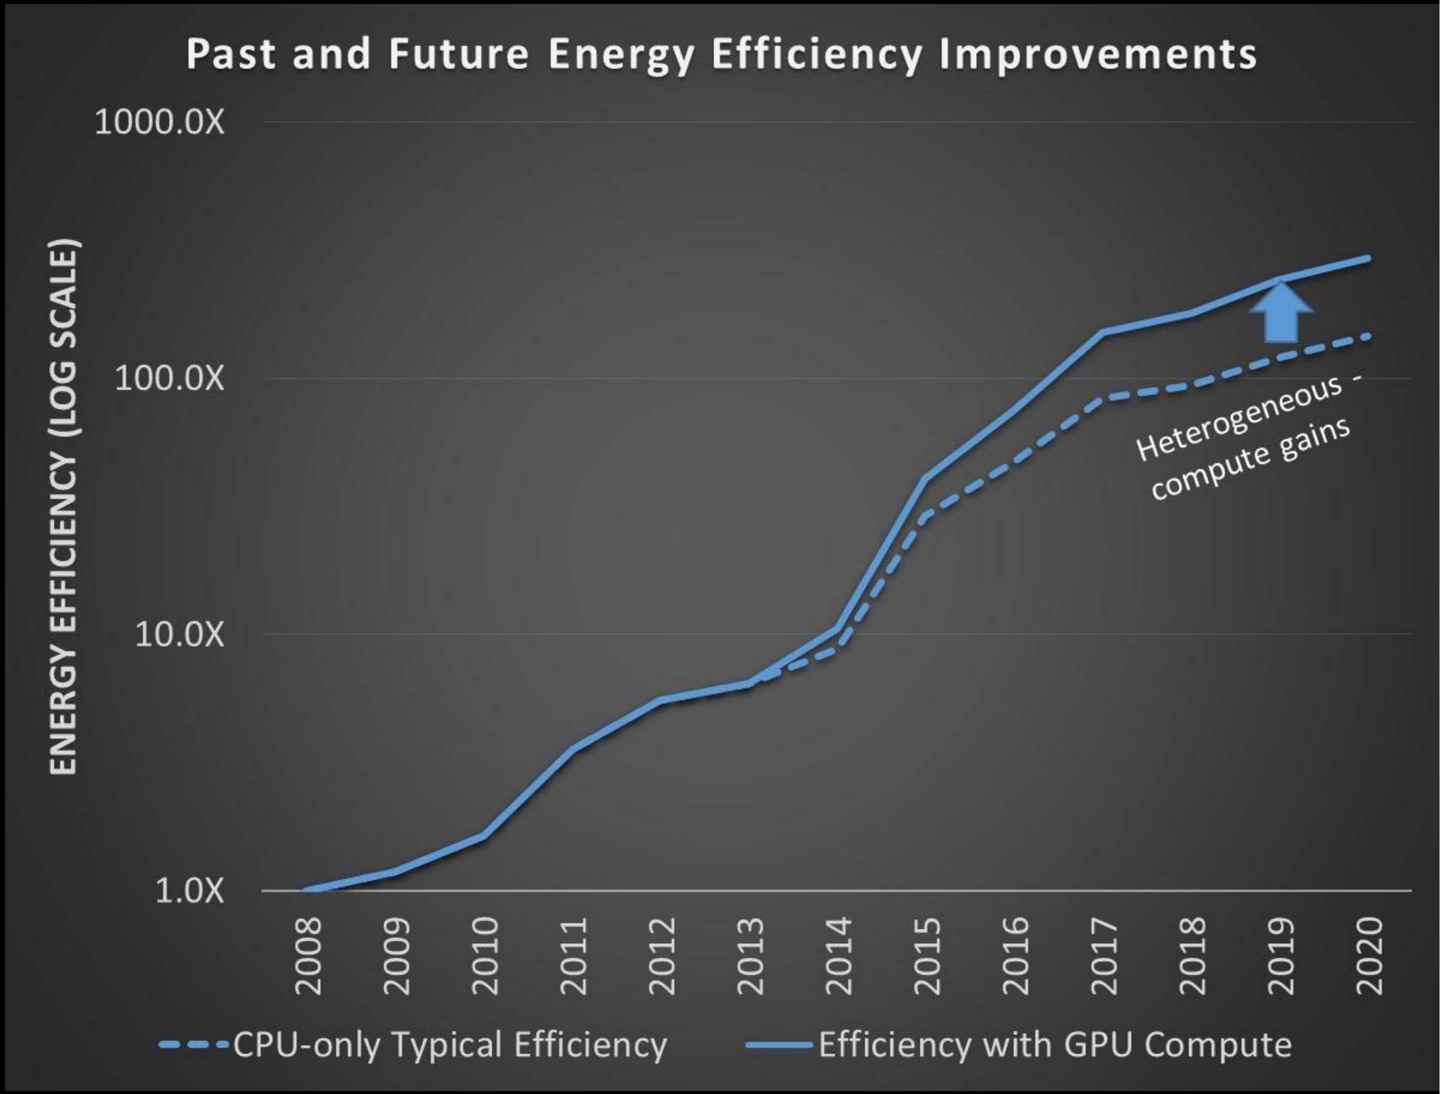

[1]Typical-use Energy Efficiency as defined by taking the ratio of compute capability as measured by common performance measures such as SpecIntRate, PassMark and PCMark, divided by typical energy use as defined by E_{TEC} (Typical Energy Consumption for notebook computers) as specified in Energy Star Program Requirements Rev 6.0 10/2013

Slide Sources

- “AMD Product and Tech Roadmaps 5.5.14”
- “Energy Efficiency Messaging” May 9, 2014
- “AMD Server Roadmap”, November 2013

DISCLAIMER & ATTRIBUTION

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors.

The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.

AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION.

AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY DIRECT, INDIRECT, SPECIAL OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

ATTRIBUTION