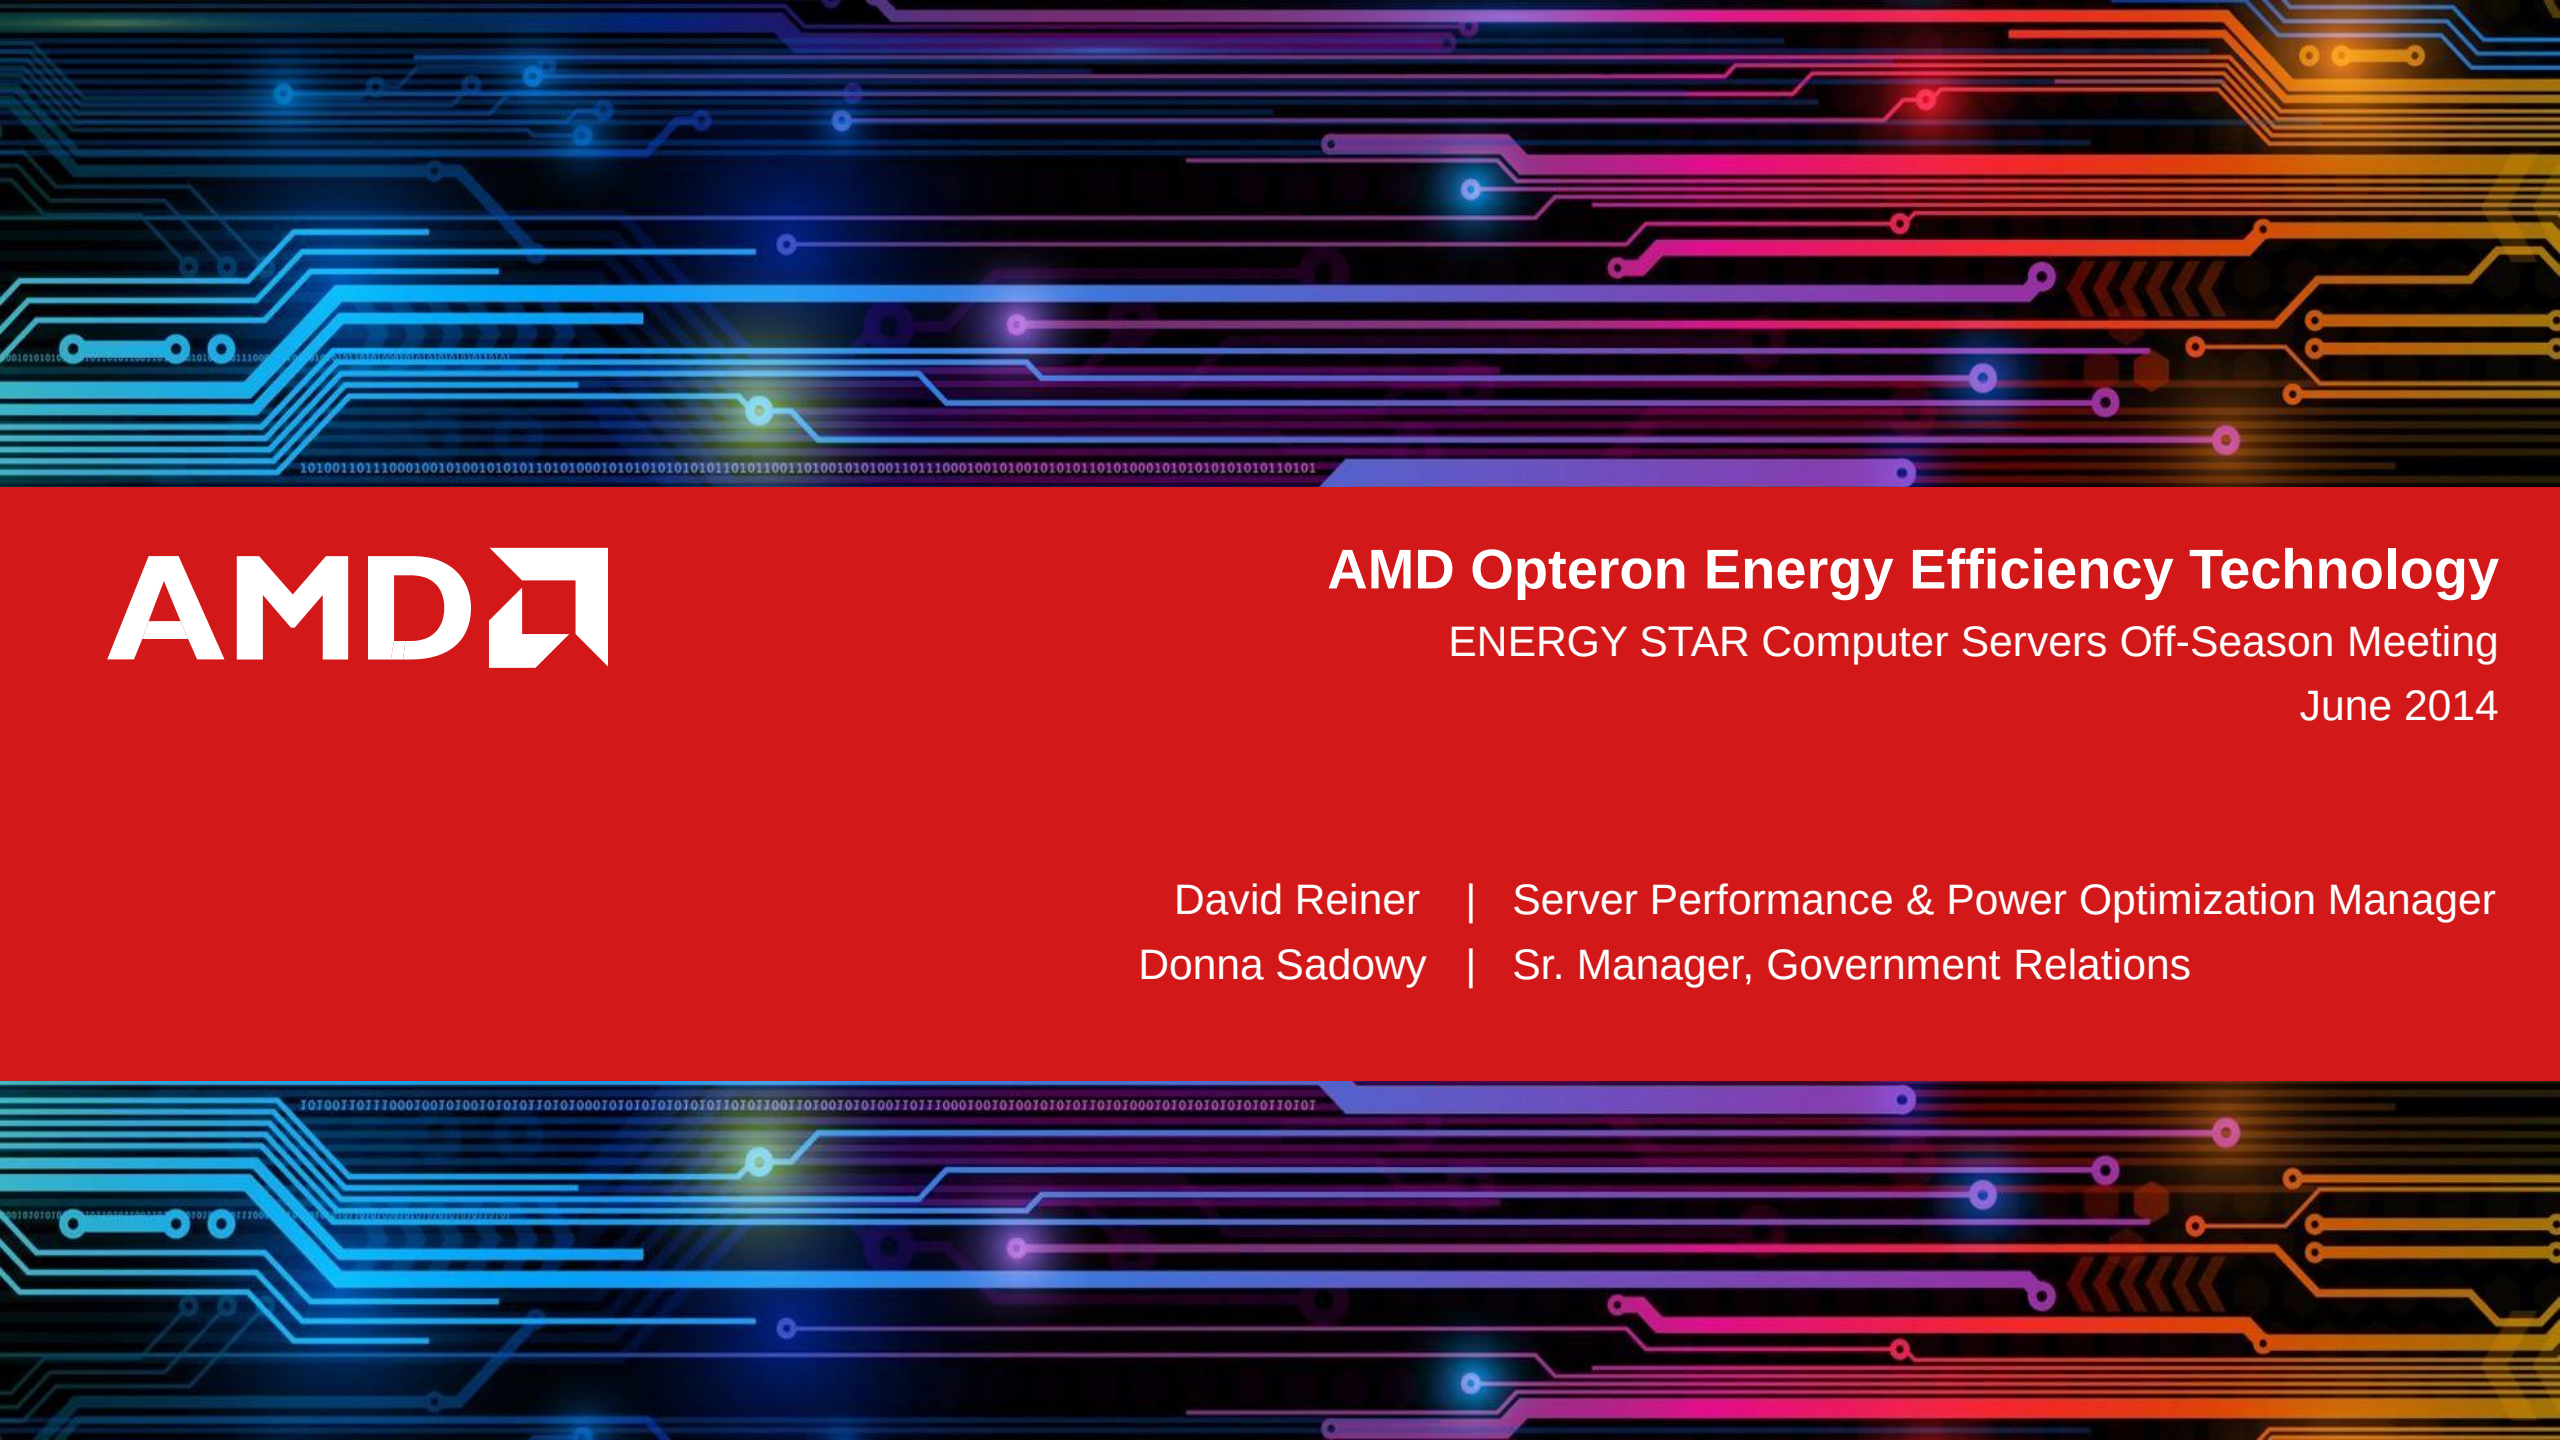

AMD Opteron Energy Efficiency Technology

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 **AMD Opteron™ X1150 CPU and X2150 APU** 4 “Jaguar” CPU Cores GCN Graphics Compute Units (APU) 9W-22W 28nm SoC	**“Berlin” CPU/APU** 4 “Steamroller” CPU Cores GCN Graphics Compute Units (APU) HSA Features (APU) 28nm **“Seattle” CPU** ARM “A57” CPU Cores 28nm SoC

A History of Energy Efficiency

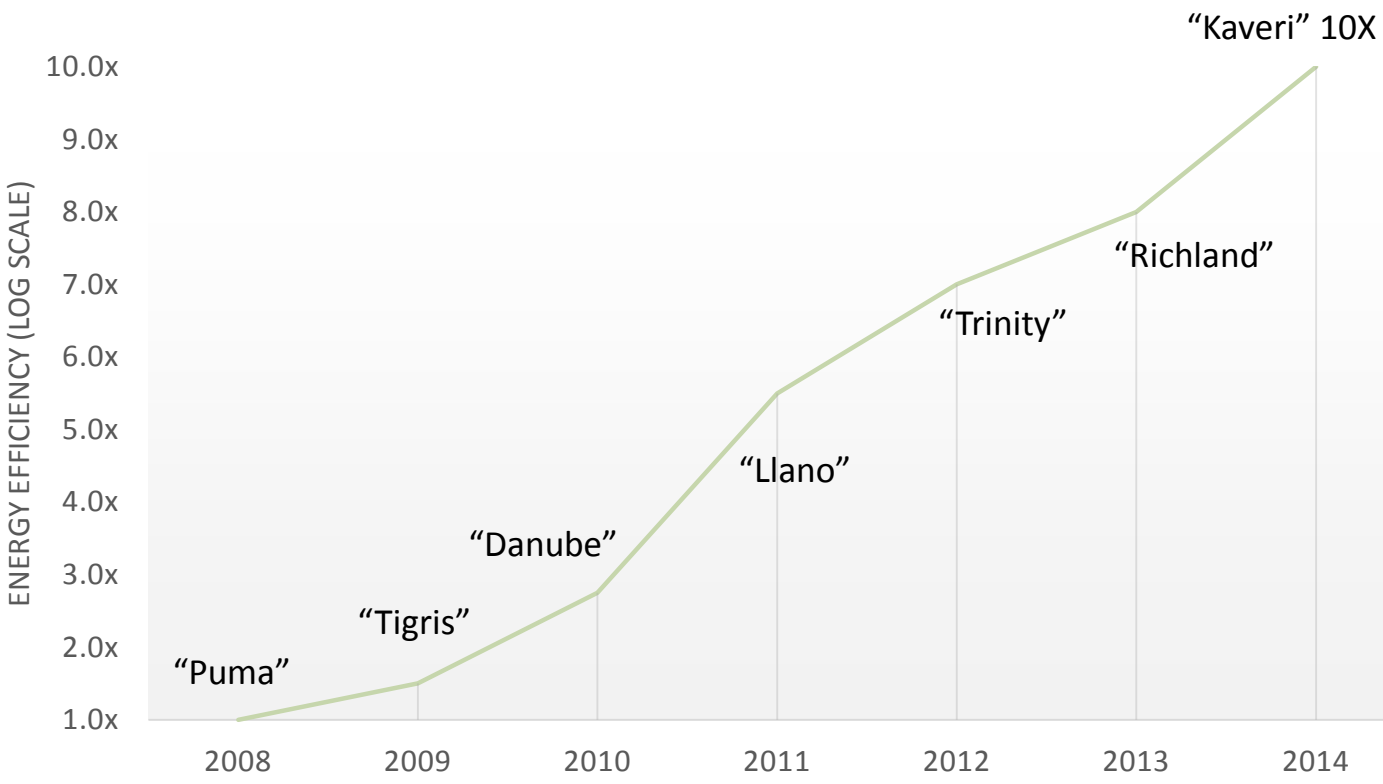

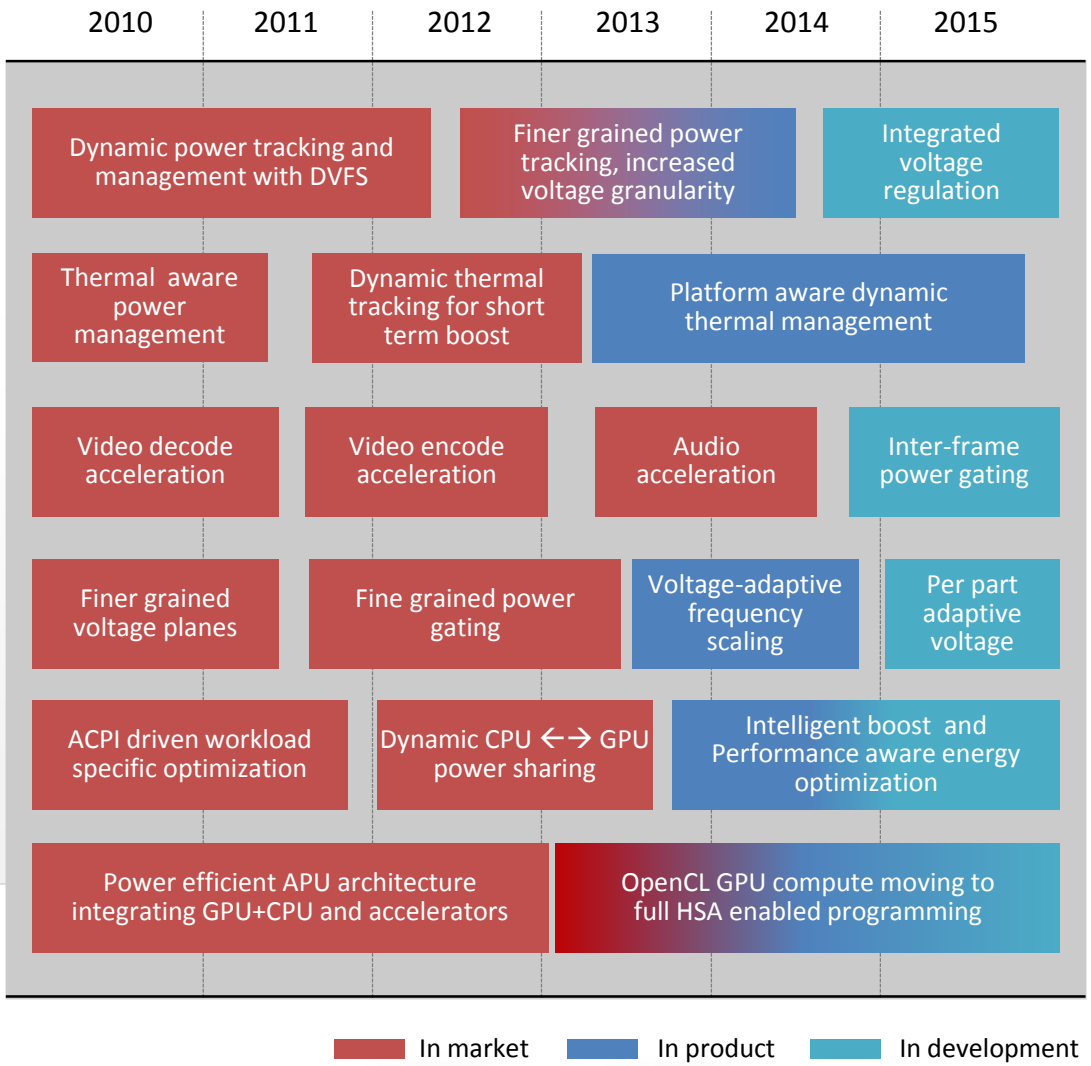

CHIP-LEVEL POWER DISTRIBUTIONS

- Power consumption (and hence performance) is set by the cooling capabilities of the platform
- Power varies a lot by workload
- We measure and manage the power of each component on the chip to generate the best performance/watt

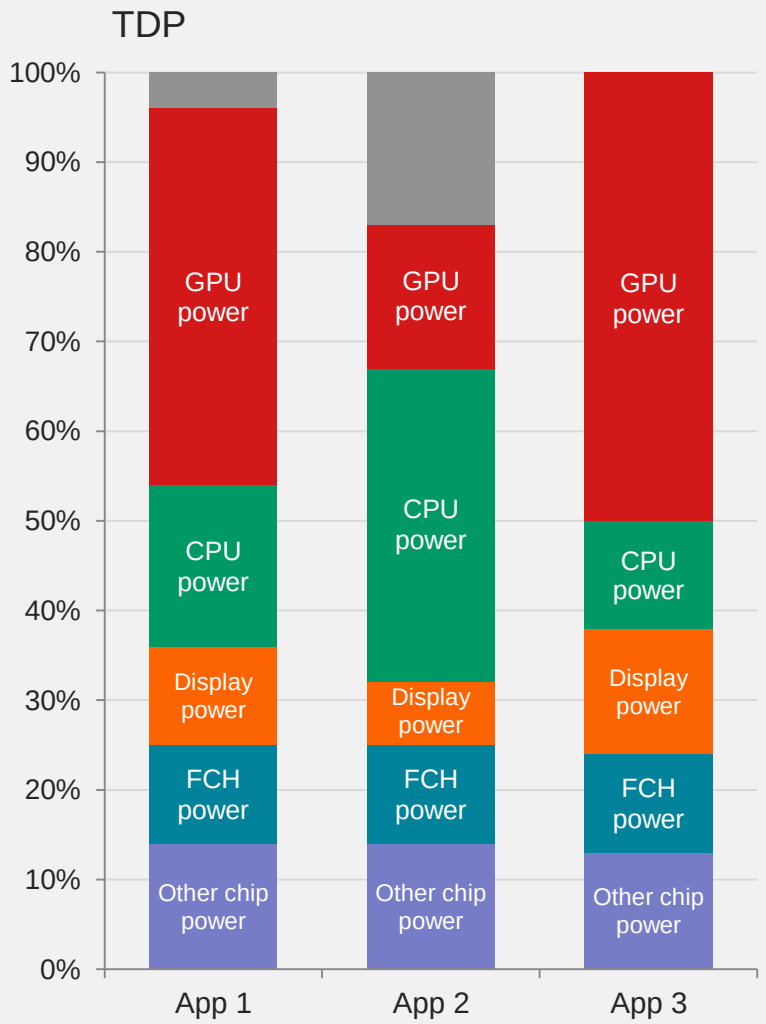

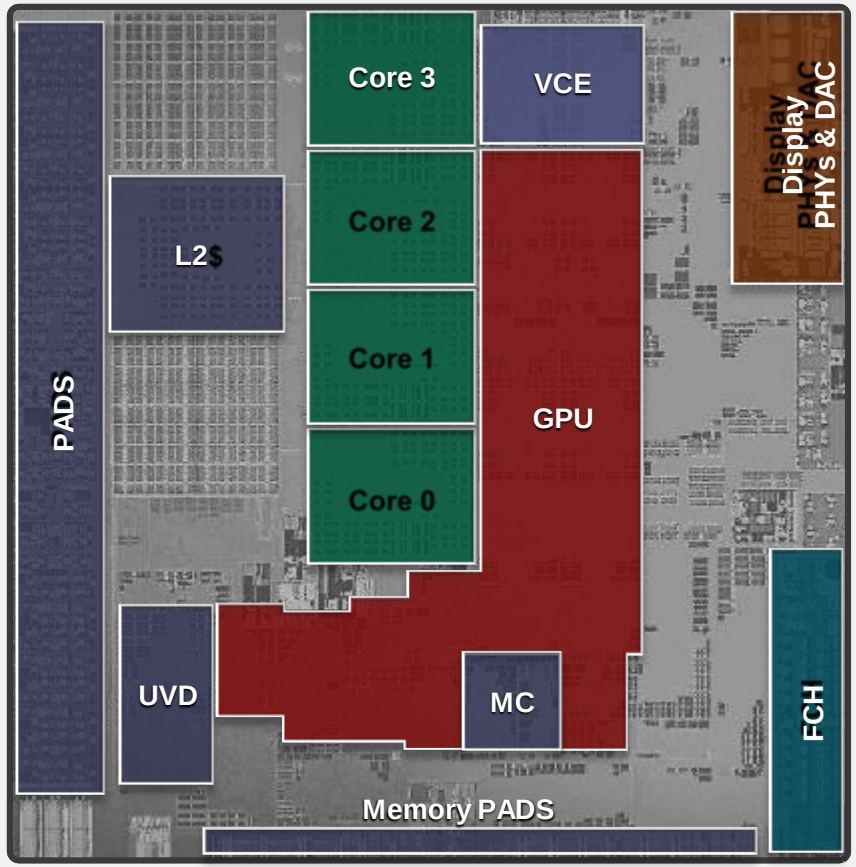

AMD Turbo Core Technology

- To manage temperature and send the power wherever it's needed, we use power monitors in all chip components
- "Kyoto" has power monitors in each CPU, the GPU, the display interface, and the FCH
- The central controller uses this information to optimize performance within thermal constraints

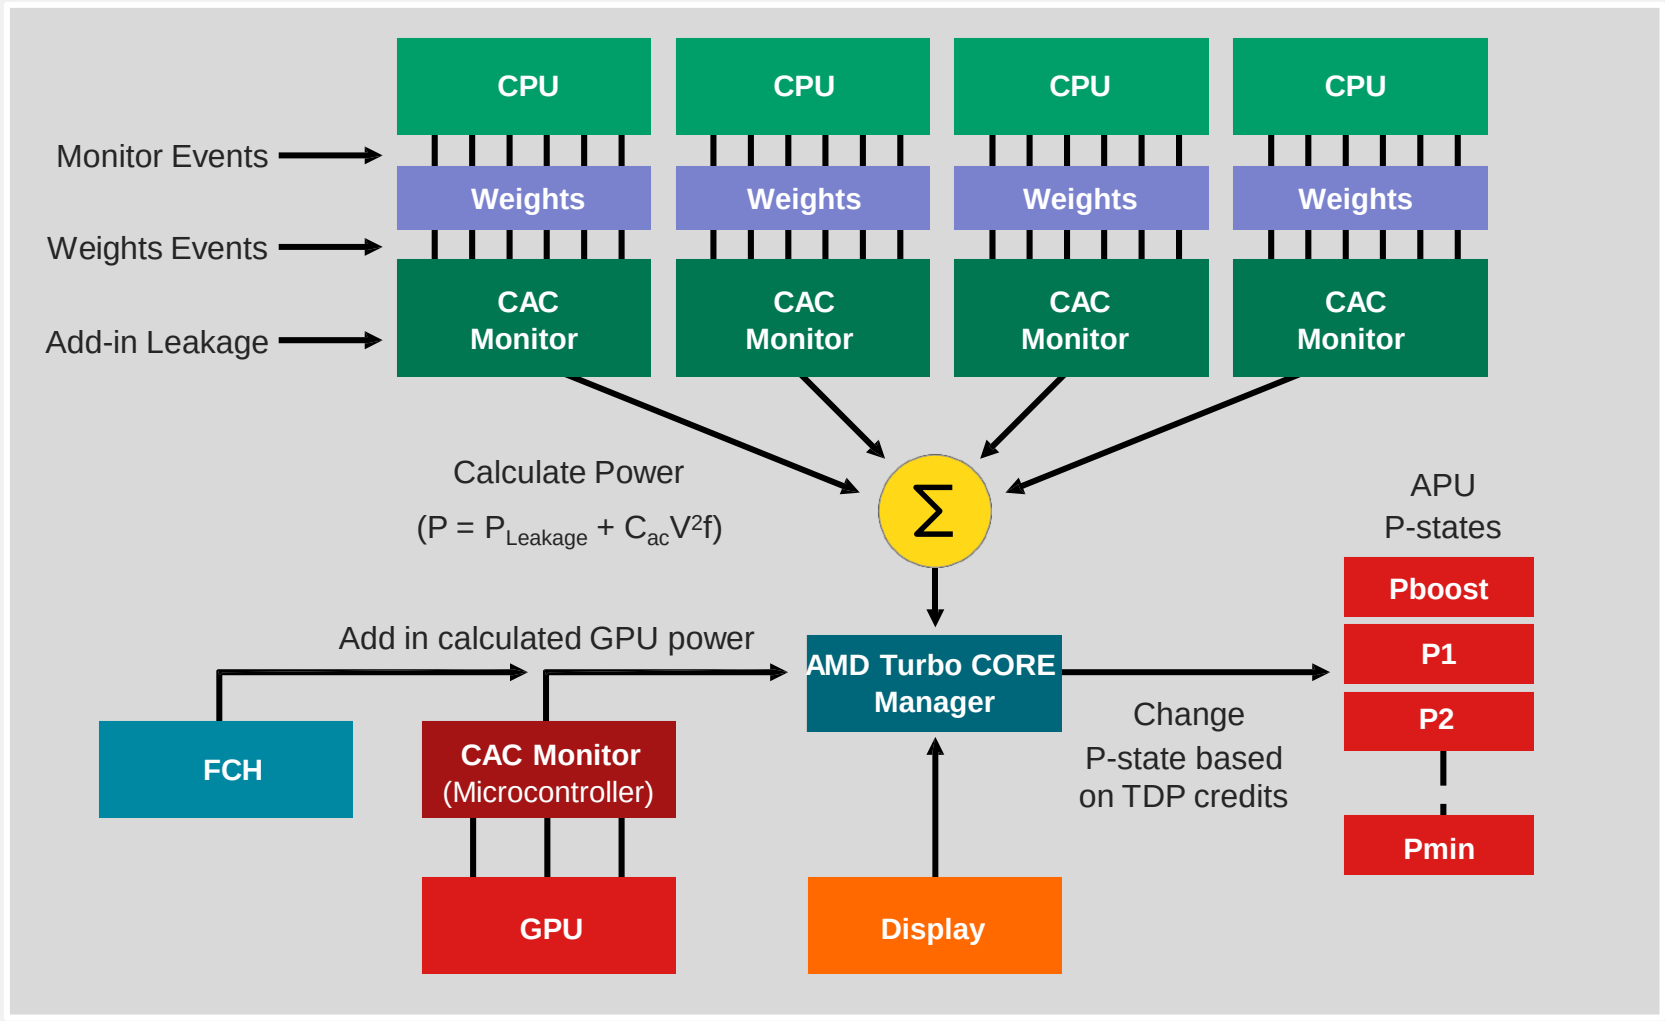

EVOLUTION OF AMD TURBO CORE TECHNOLOGY

Year	Processor	Boosting decision based on	Notes
2010	**AMD Phenom™ II**	▪ **Number of cores active**	▪ **Single boost Pstate** used if half or more cores are inactive ▪ Coarse-grain power margin exploited
2011	**1st-Generation AMD A-Series APU**	▪ **Calculated power**	▪ **Unidirectional power transfer** between thermal entities ▪ GPU→CPU ▪ **Exploit fine-grain power margin**

2012	**2nd-Generation AMD A-Series APU**	▪ Calculated power
		▪ **Calculated temperature**
			▪ **Bidirectional power transfer** between thermal entities
			  ▪ GPU→CPU
			  ▪ CPU→GPU
			▪ **Exploit temperature margin**

2013	**3rd-Generation AMD A-Series APU ("Richland“)**	▪ Calculated power ▪ Calculated temperature ▪ **Measured/Sensor temperature** ▪ **Efficiency of power usage** by individual entities (CPU, GPU, etc.)	▪ **Designed to more effectively exploit temperature margin** by detecting favorable thermal conditions in real time ▪ **Intelligent Boost**

CHIP-LEVEL POWER DISTRIBUTIONS: GPU-CENTRIC

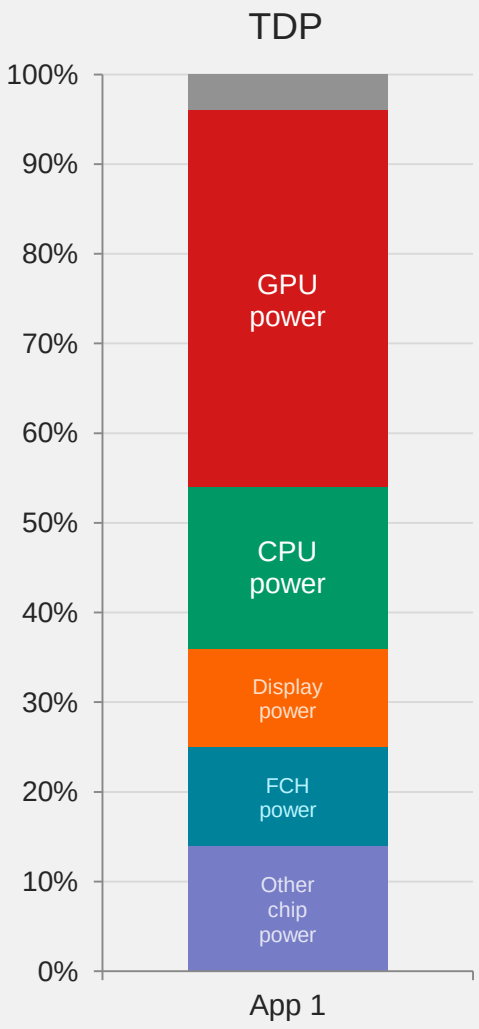

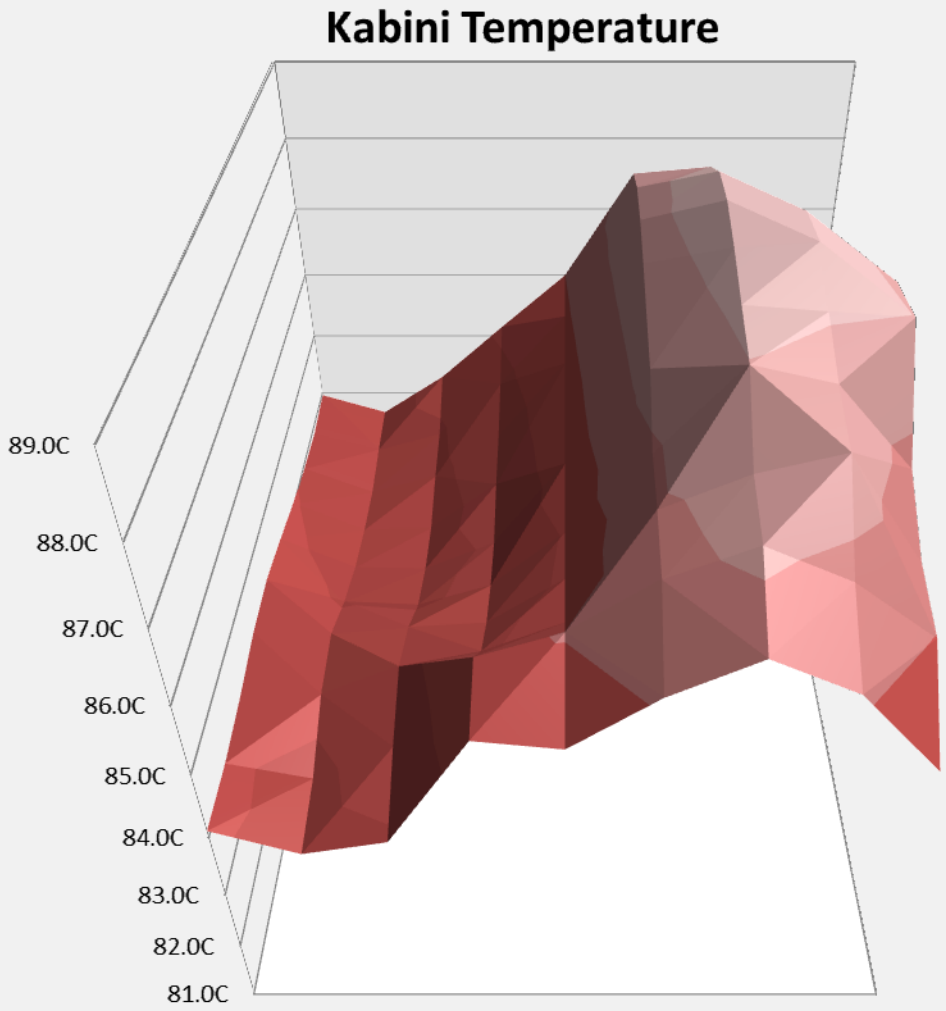

Lower-power cores serve as a heat sink for the active GPU

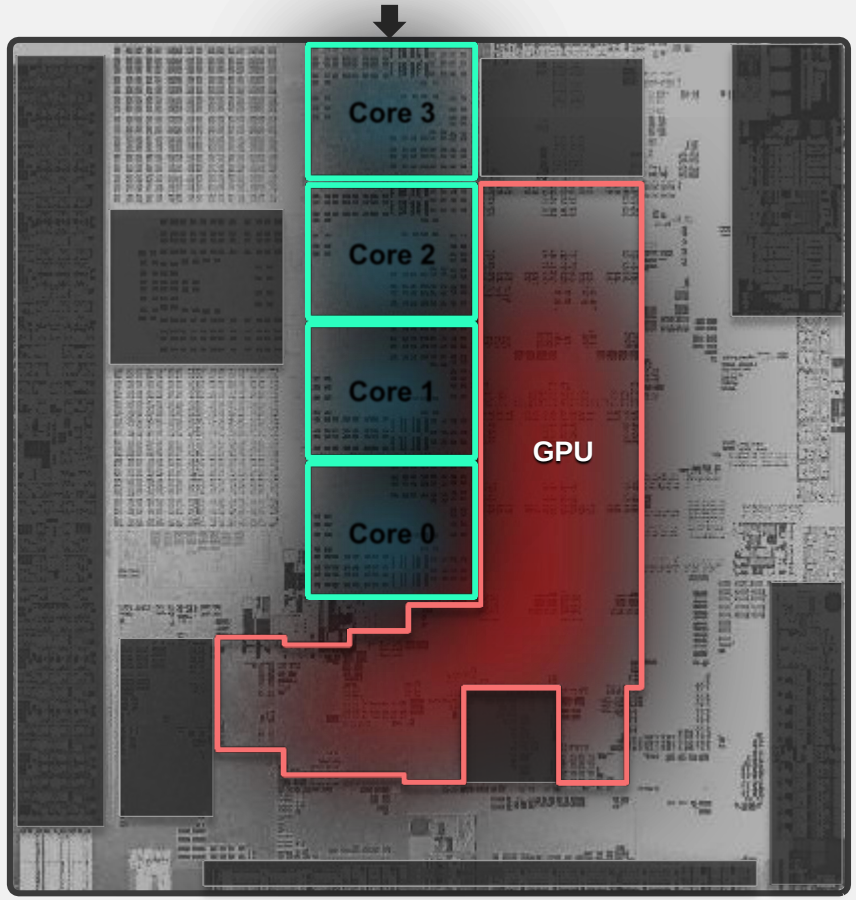

CHIP-LEVEL POWER DISTRIBUTIONS: CPU-CENTRIC

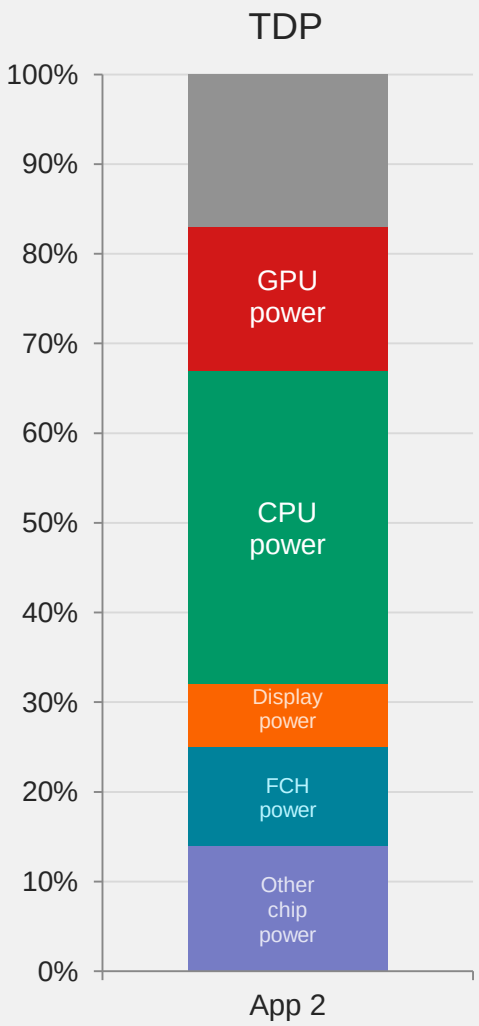

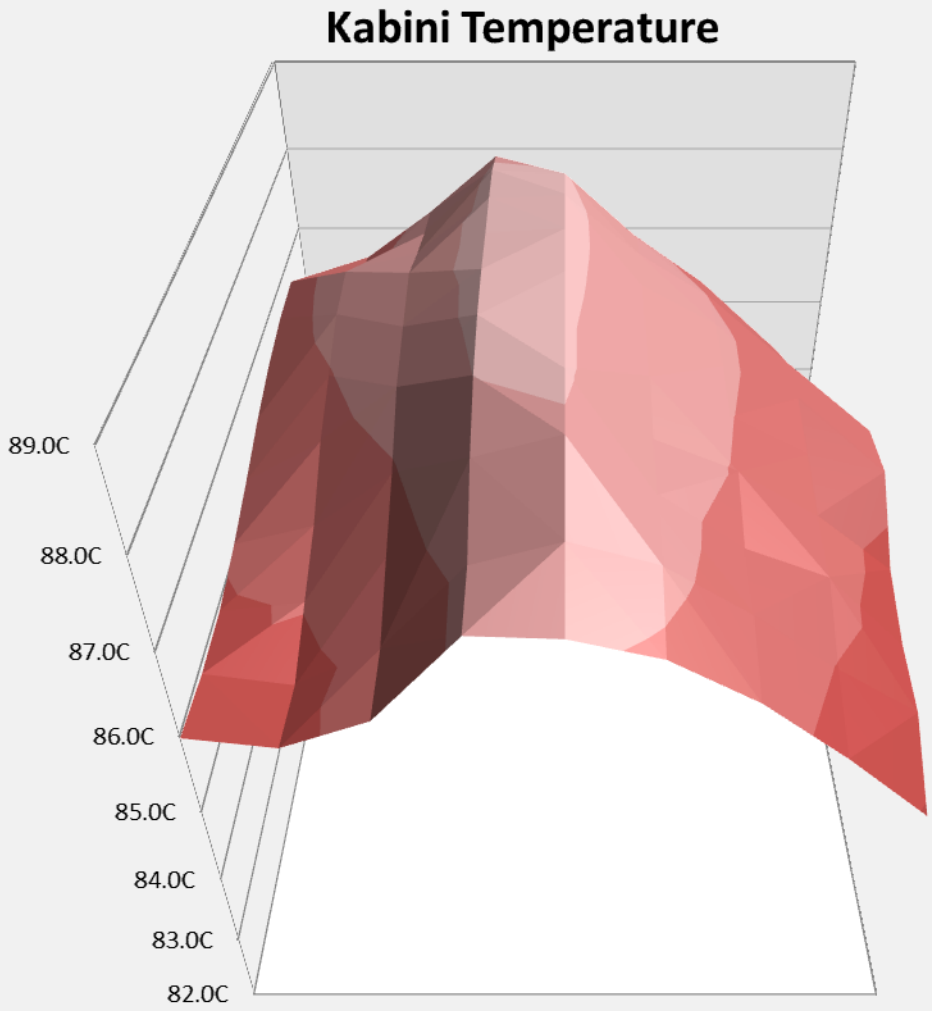

Lower-power GPU serves as a heat sink for the active CPUs

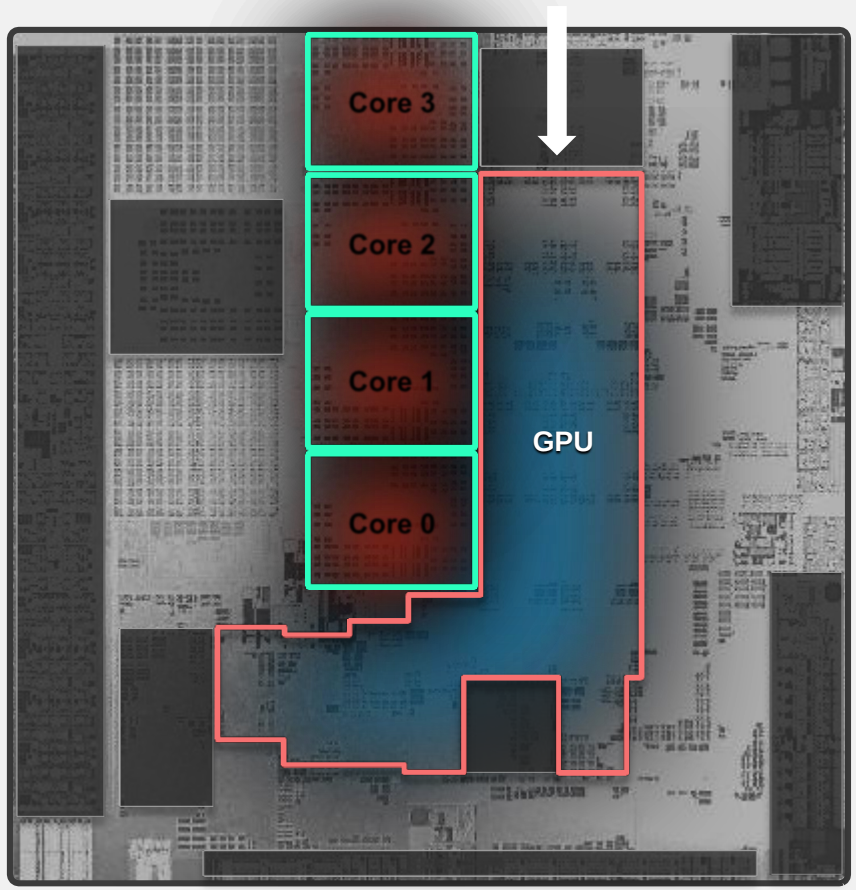

Avoiding power waste with Intelligent Boost control

- Intelligent Boost is designed to avoid power waste that results from boosting applications that benefit very little from higher frequency
- Enables long battery life and cool operation while maintaining great performance
- Power management micro-controller tracks application behavior real-time to determine frequency sensitivity
- Boost behavior is adjusted accordingly

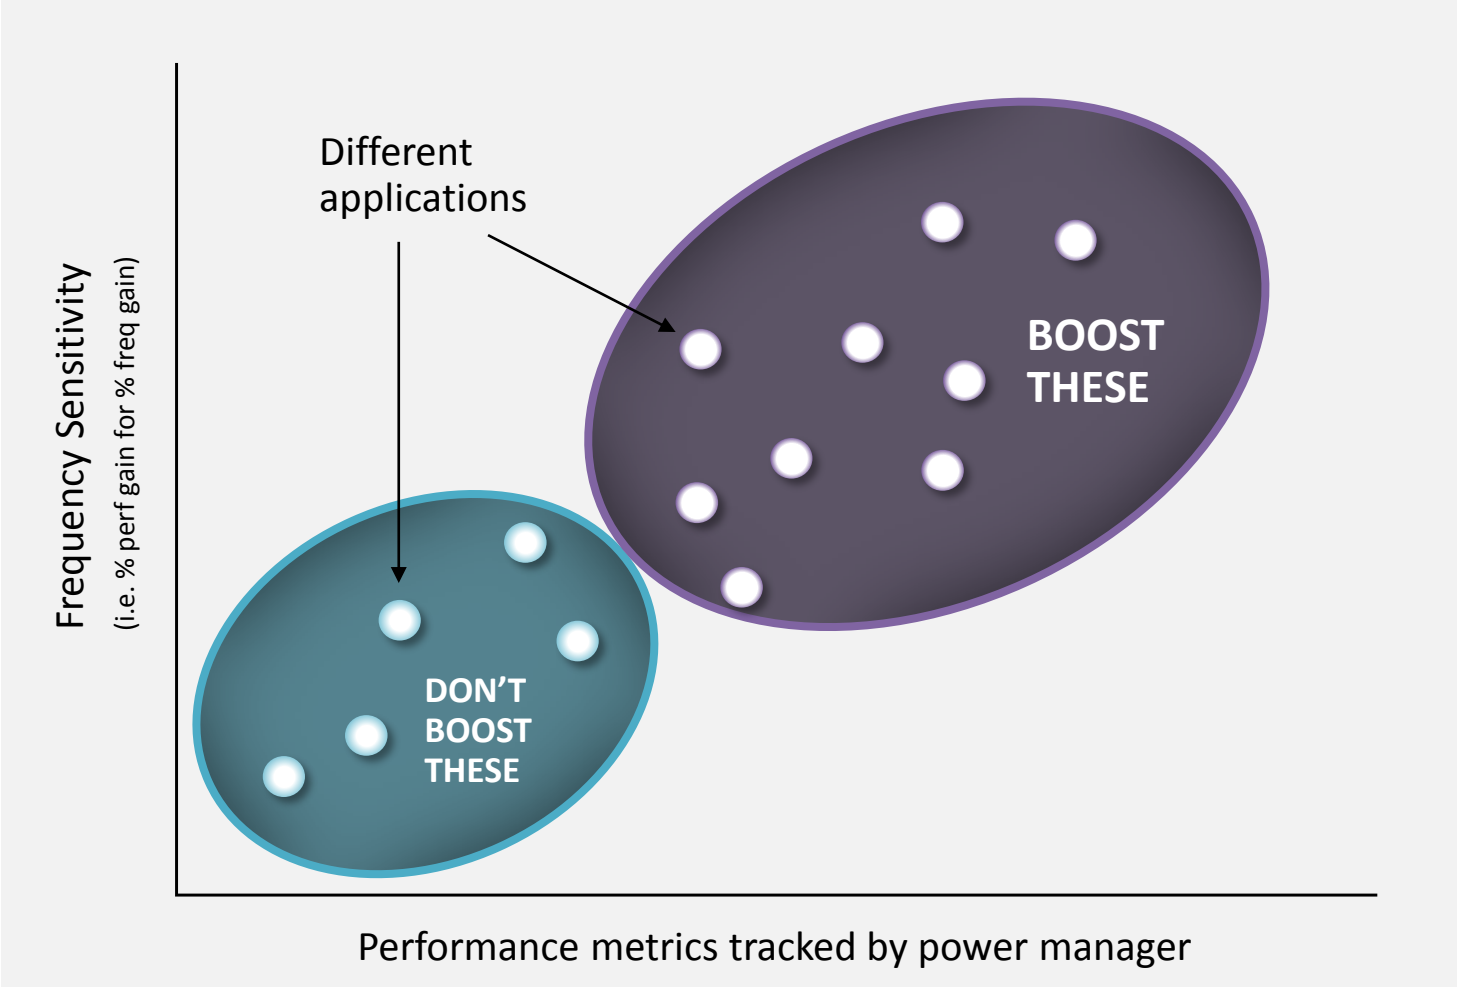

FABRICS REDUCE THE POWER CONSUMPTION OF EVERY SERVER

- ~70% of energy consumption in servers from components beyond CPU
- Eliminate unnecessary components and functions
- Remove tiers of networking equipment and thousands of cables

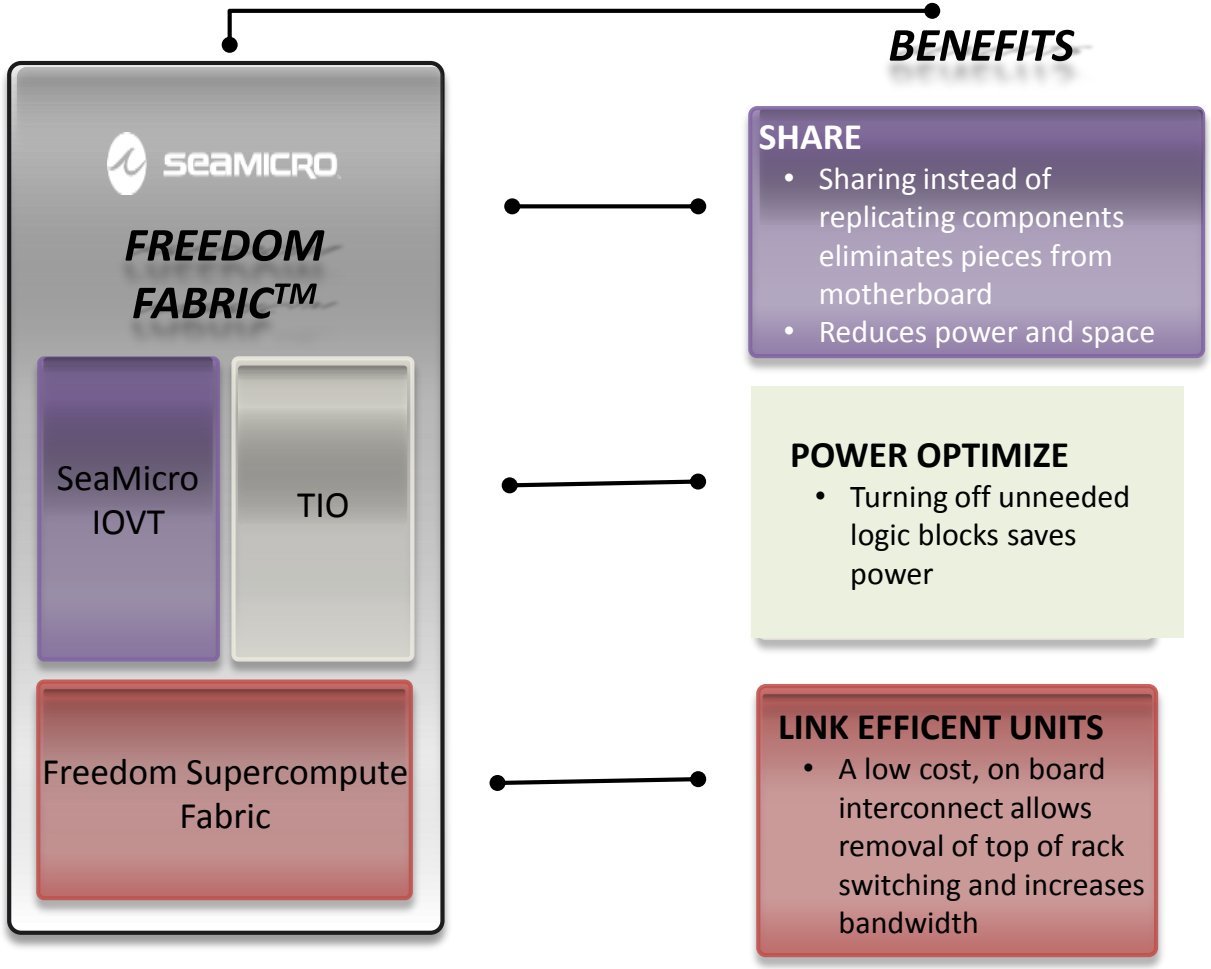

Slide Sources

- ““Richland” Client APU” Presentation by Praveen Dongara, Lloyd Bircher, John Darrilek - Hot Chips 25, August 2013
- “AMD Product and Tech Roadmaps 5.5.14”
- “Energy Efficiency” by Sam Naffziger, June 16, 2014
- “AMD “Kabini” APU SOC” by Dan Bouvier, Ben Bates, Walter Fry, Sreekanth Godey – Hot Chips 25, August 2013
- “Energy Efficiency Messaging” May 9, 2014
- “AMD Advanced Power Management” by Sam Naffziger, April 2014

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 WILLAMD 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