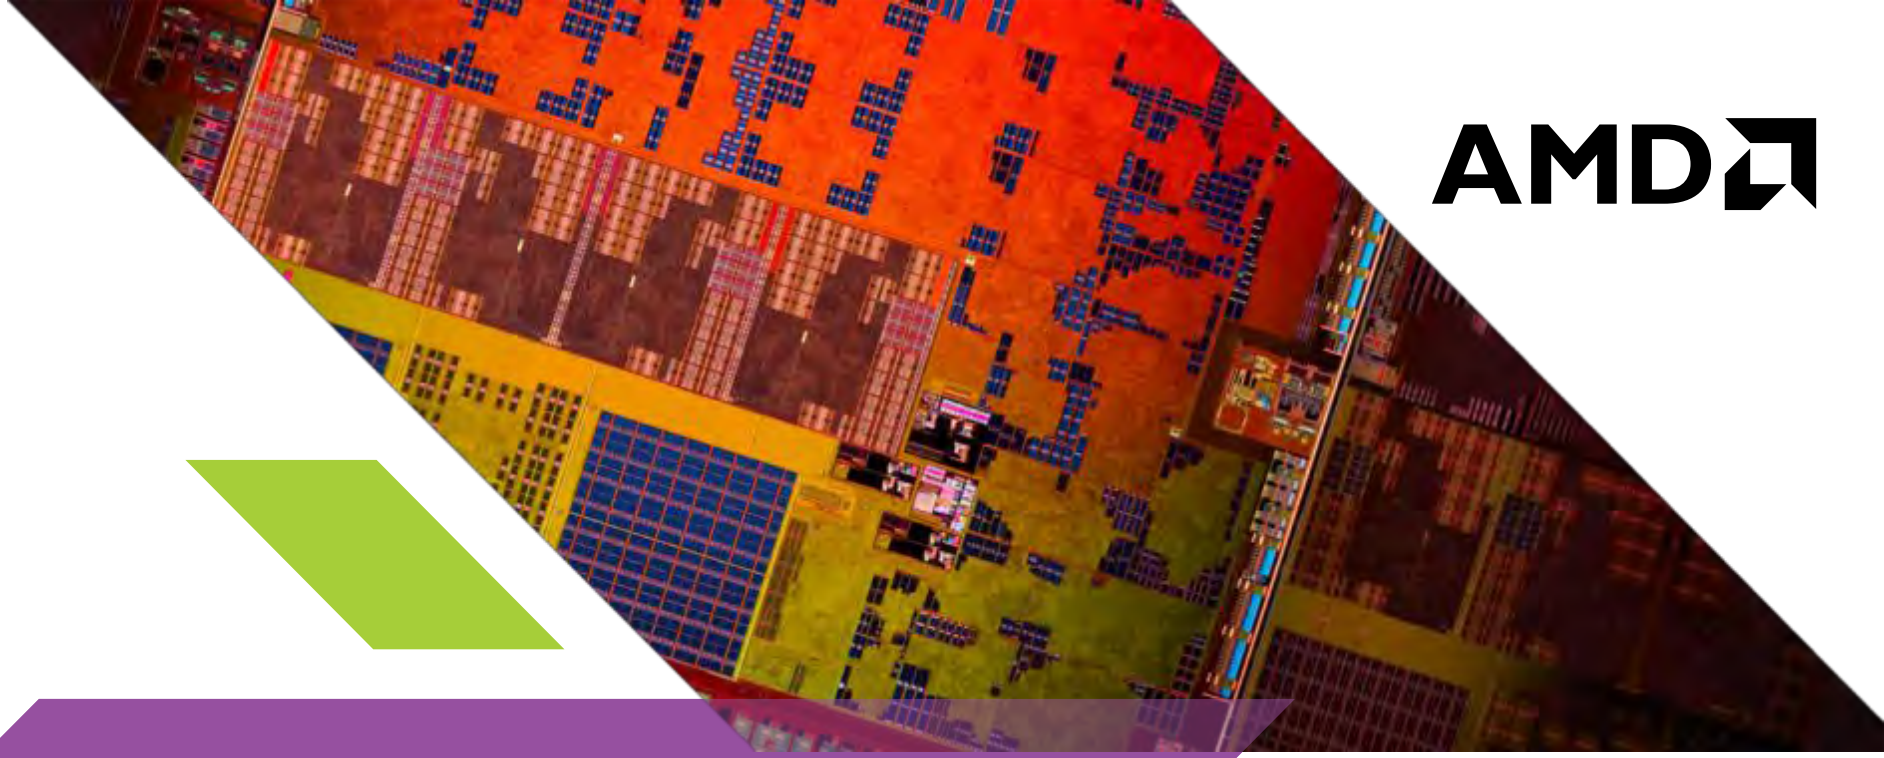

AMD POWER MANAGEMENT

SRILATHA MANNE

11/14/2013

OUTLINE

- APUs and HSA
- Richland APU
- APU Power Management
- Monitoring and Management Features

APU: ACCELERATED PROCESSING UNIT

- The APU has arrived and it is a great advance over previous platforms
- Combines scalar processing on CPU with parallel processing on the GPU and high-bandwidth access to memory
- Advantages over other forms of compute offload:
  - Easier to program
  - Easier to optimize
  - Easier to load-balance
  - Higher performance
  - Lower power

© Copyright 2012 HSA Foundation. All Rights Reserved.

WHAT IS HSA?

AMD

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*

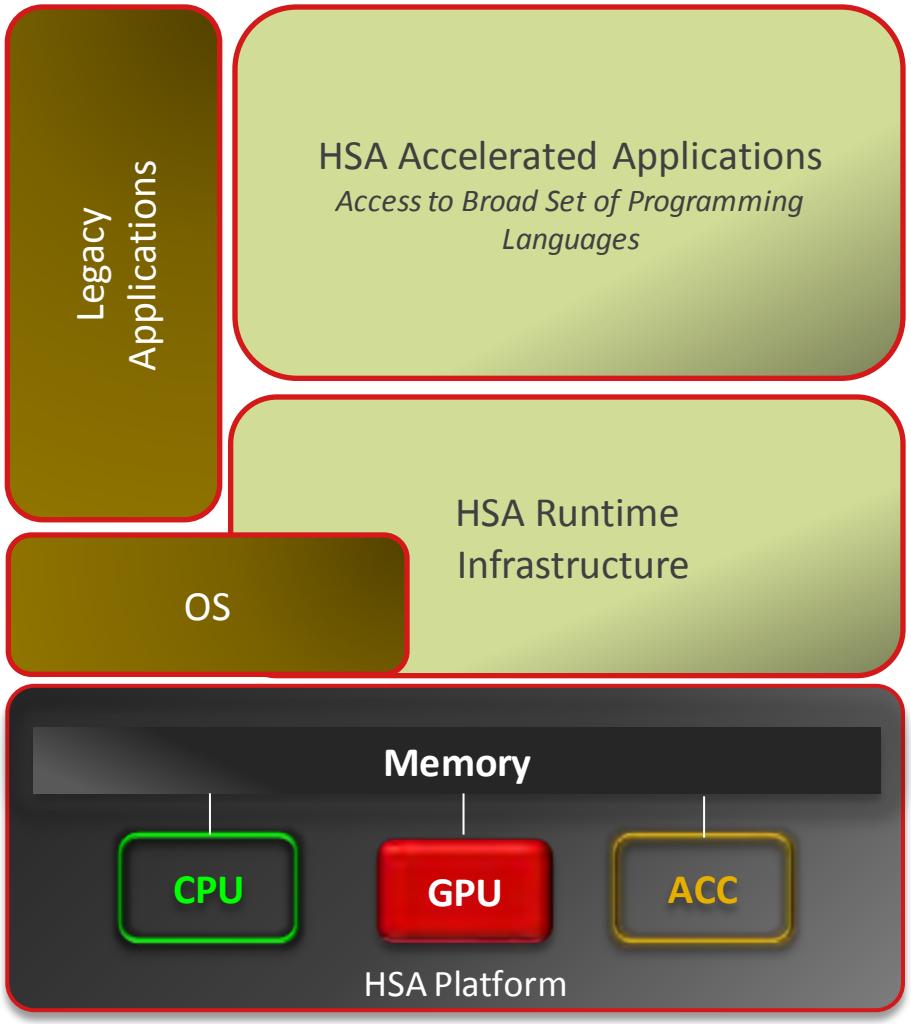

A NEW ERA OF PROCESSOR PERFORMANCE

Single-core Era

Enabled by:
- ✓ Moore's Law
- ✓ Voltage Scaling

Constrained by:
- ✗ Power
- ✗ Complexity

Assembly ➡ C/C++ ➡ Java …

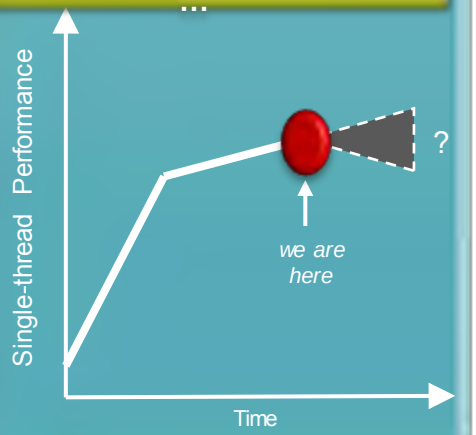

Multi-core Era

Enabled by:
- ✓ Moore's Law
- ✓ SMP architecture

Constrained by:
- ✗ Power
- ✗ Parallel SW
- ✗ Scalability

pthreads ➡ OpenMP / TBB …

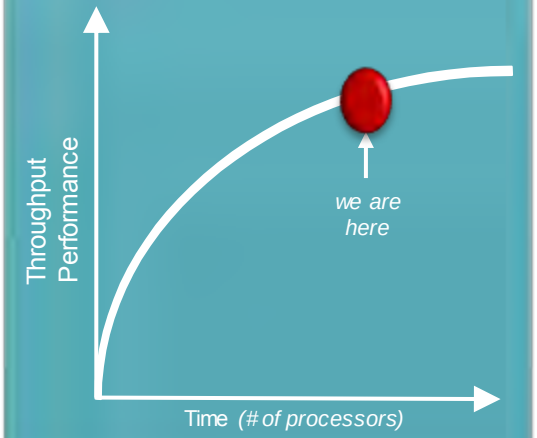

Heterogeneous Systems Era

Enabled by:
- ✓ Abundant data parallelism
- ✓ Power efficient GPUs

***Temporarily* Constrained by:**
- ✗ Programming models
- ✗ Comm.overhead

Shader ➡ CUDA ➡ OpenCL ➡ C++ and Java

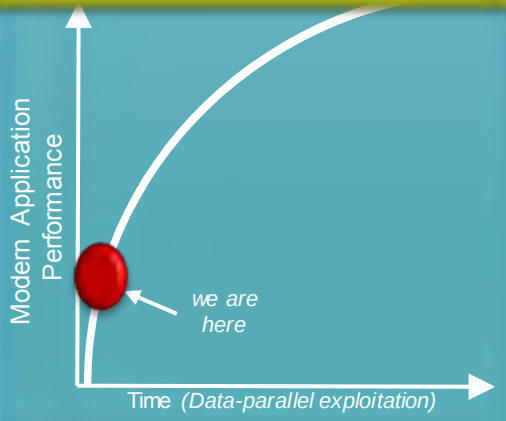

© Copyright 2012 HSA Foundation. All Rights Reserved.

CASCADE DEPTH ANALYSIS

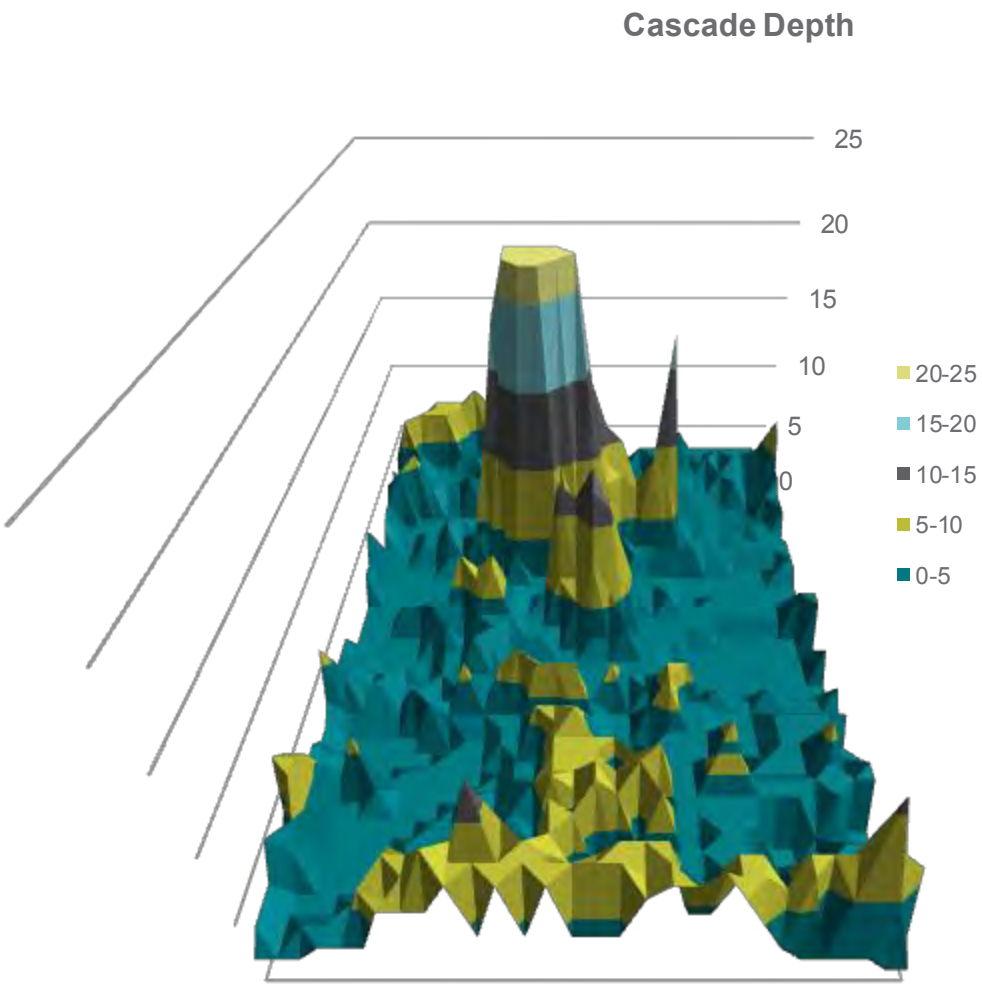

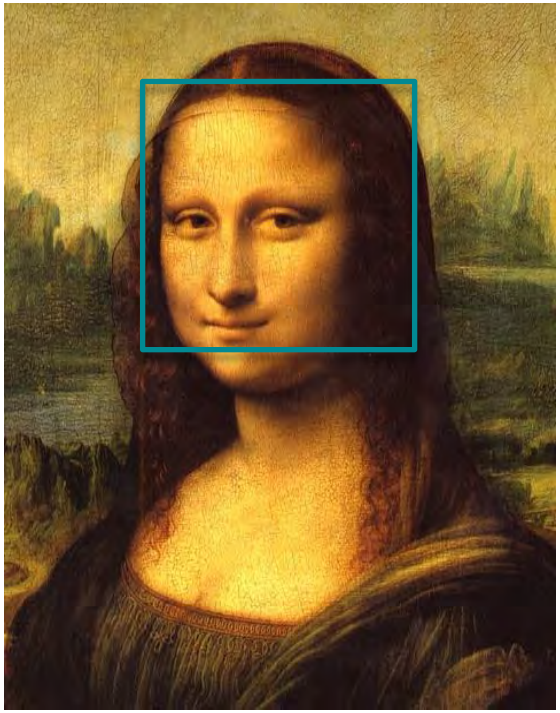

© Copyright 2012 HSA Foundation. All Rights Reserved.

PROCESSING TIME/STAGE

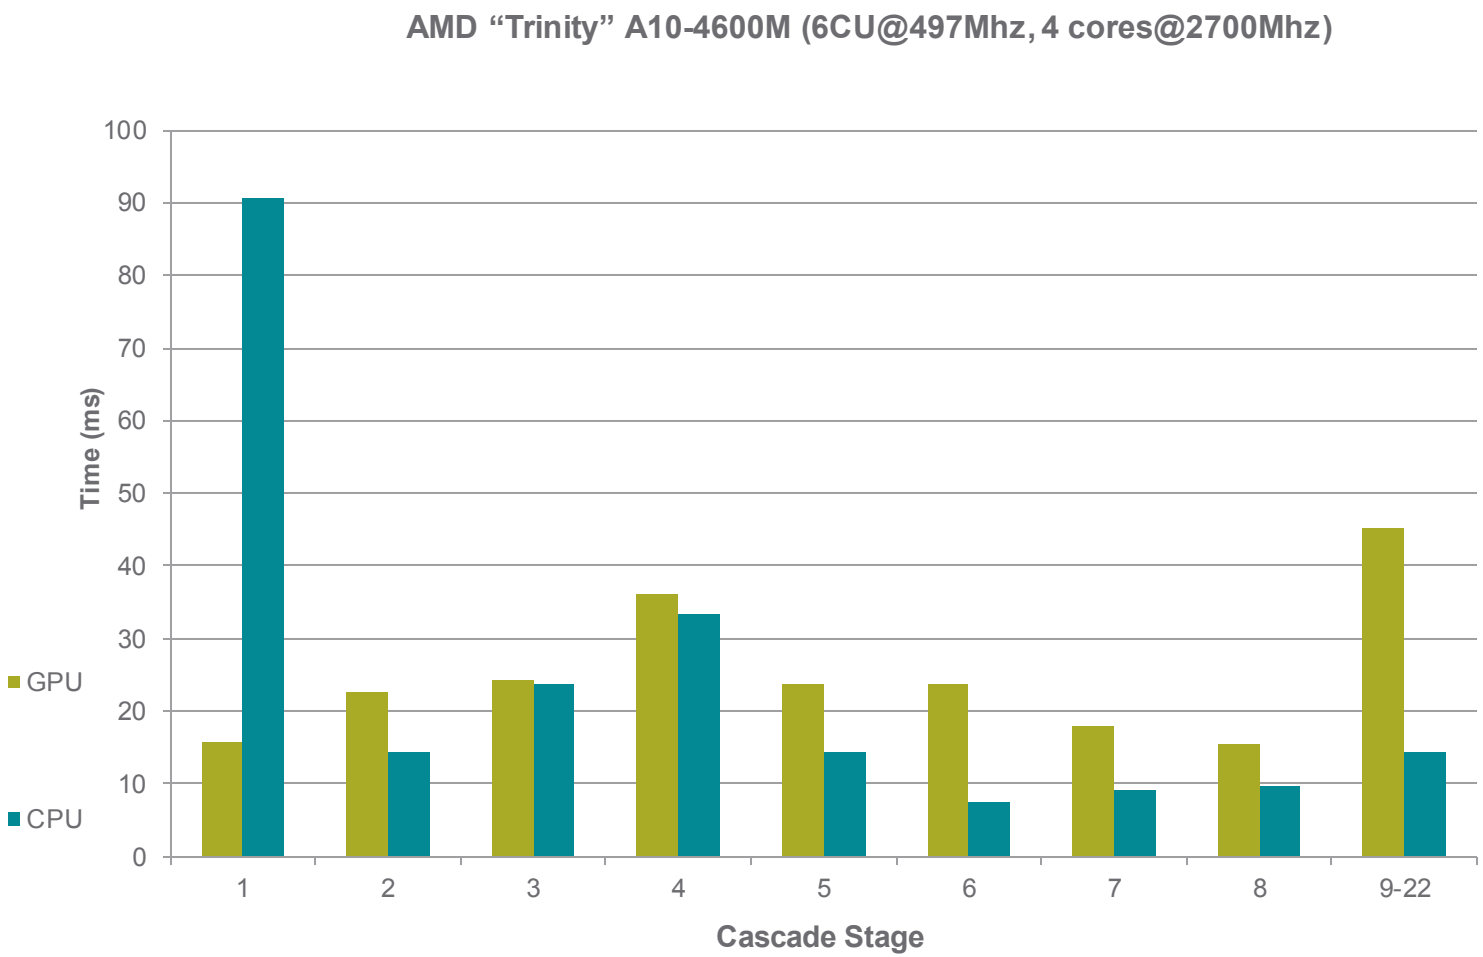

AMD A10 4600M APU with Radeon™ HD Graphics; CPU: 4 cores @ 2.3 MHz (turbo 3.2 GHz); GPU: AMD Radeon HD 7660G, 6 compute units, 685MHz; 4GB RAM; Windows 7 (64-bit); OpenCL™ 1.1 (873.1)

© Copyright 2012 HSA Foundation. All Rights Reserved.

PERFORMANCE: CPU vs. GPU

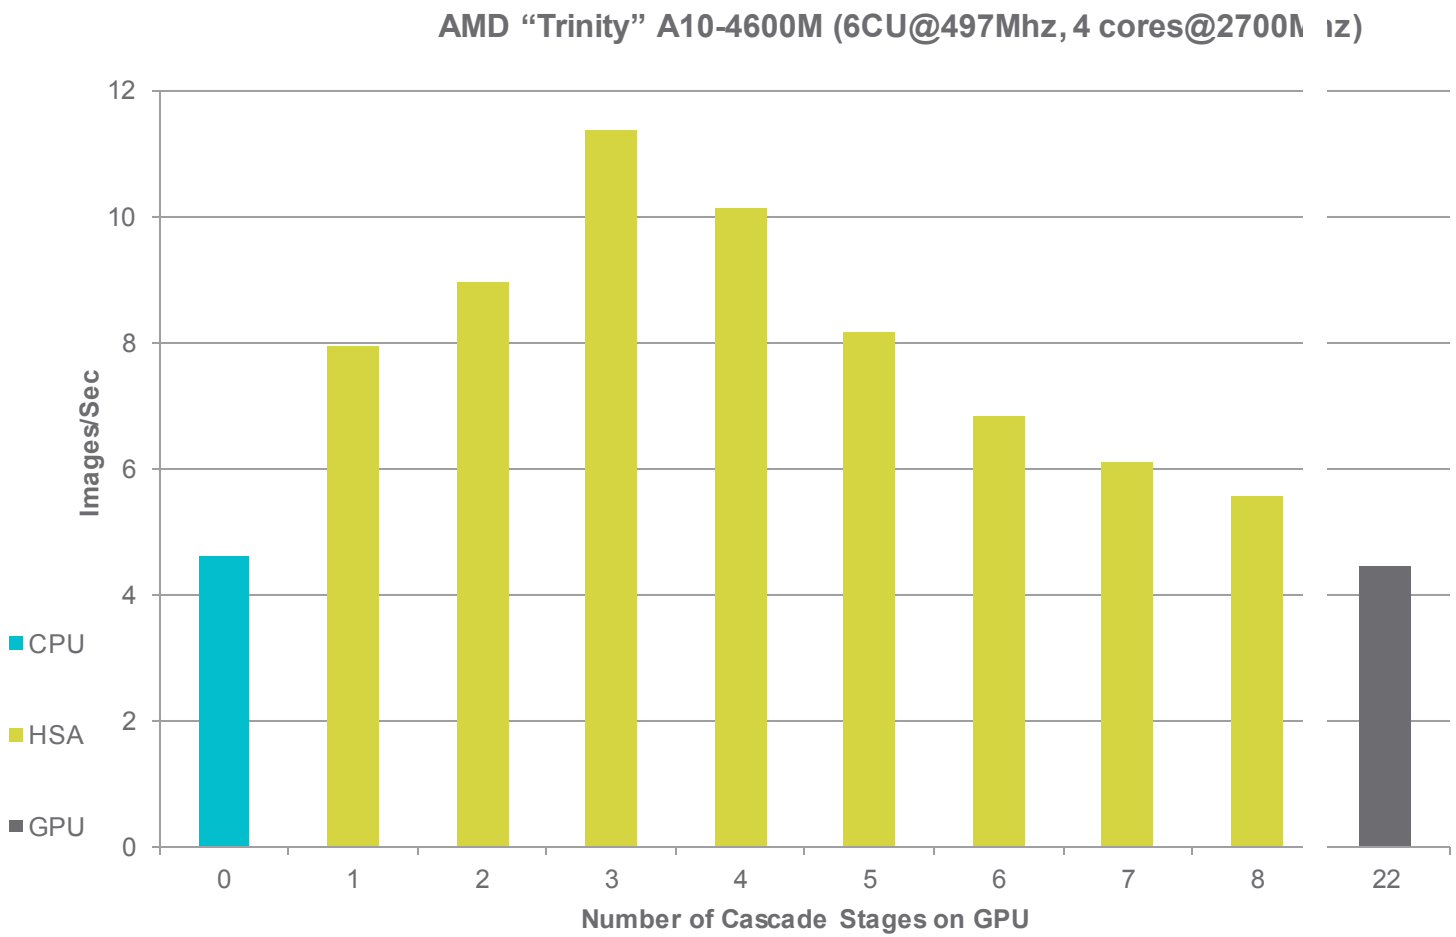

AMD A10 4600M APU with Radeon™ HD Graphics; CPU: 4 cores @ 2.3 MHz (turbo 3.2 GHz); GPU: AMD Radeon HD 7660G, 6 compute units, 685MHz; 4GB RAM; Windows 7 (64-bit); OpenCL™ 1.1 (873.1)

© Copyright 2012 HSA Foundation. All Rights Reserved.

Richland APU

APU HARDWARE

AMD "Richland" incorporates:

- Two “Piledriver” high-performance x86 modules (core-pairs)
- 2 MB shared L2 cache per x86 module
- AMD Radeon™ HD 8000 series DirectX® 11-capable GPU with six compute units
- Next-generation media acceleration technology
- Dual 64-bit memory channel supporting up to DDR3-2133
- Integrated DisplayPort 1.2 interfaces
- PCI Express® I/O Generation 2 interfaces

"Richland" is implemented in a 32-nm SOI node2+ high-K metal gate process technology

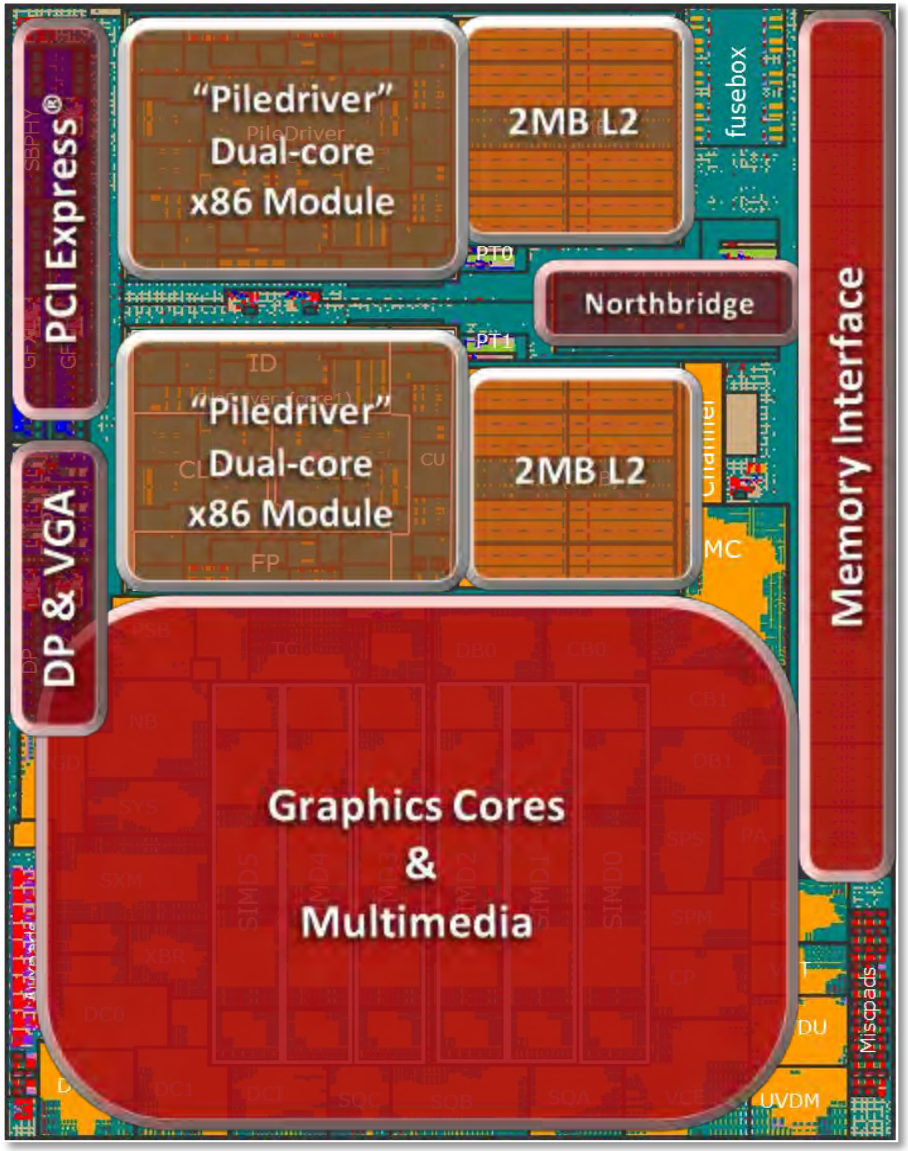

APU HARDWARE

Advantages

- Scalar and parallel capabilities
- Close physical proximity
  - Shared system memory
  - Reduced communication overhead
  - Fine-grained interaction
  - Fine-grained power management
- HSA framework
  - Flexible and powerful programming interface
  - GPU as first-class compute
  - Unified, coherent address space
  - Potential for other accelerators

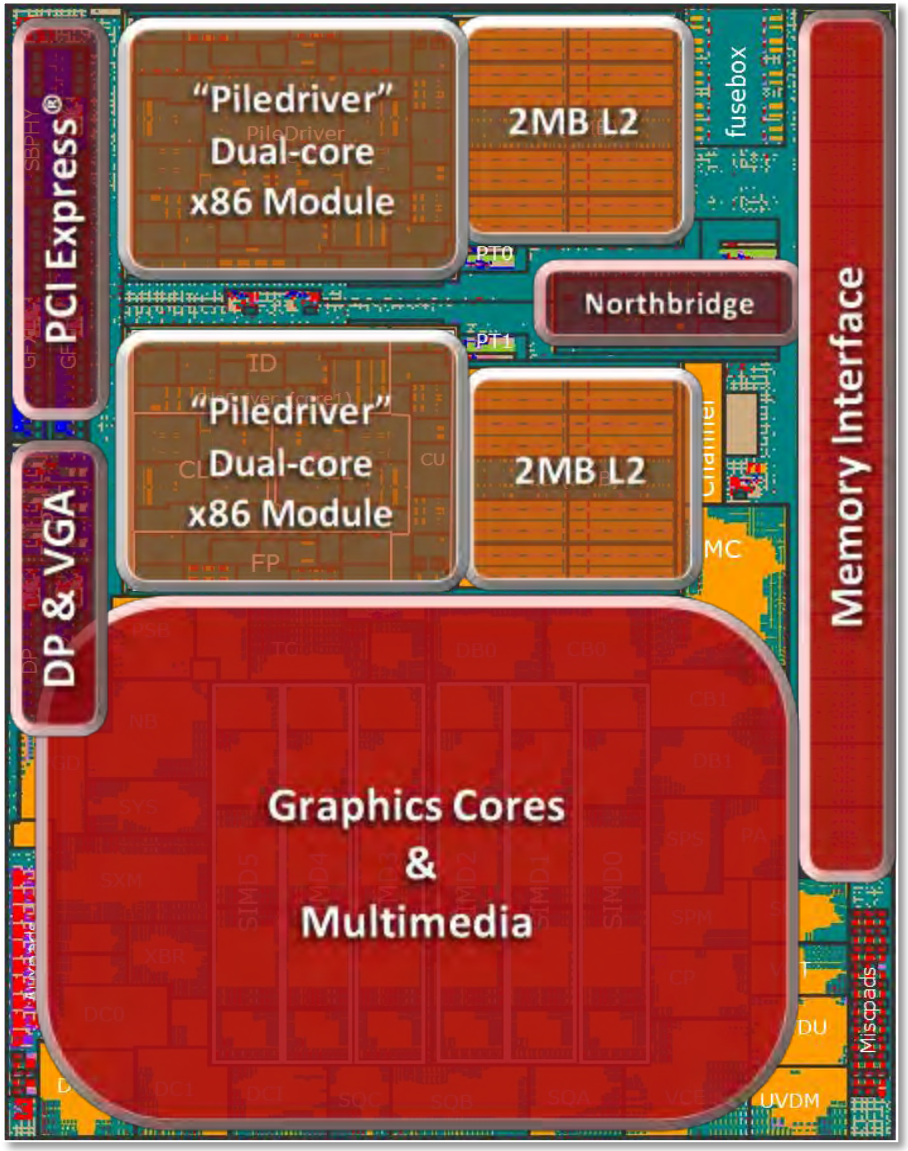

APU Power Management

AMD TURBO CORE TECHNOLOGY
MOTIVATION

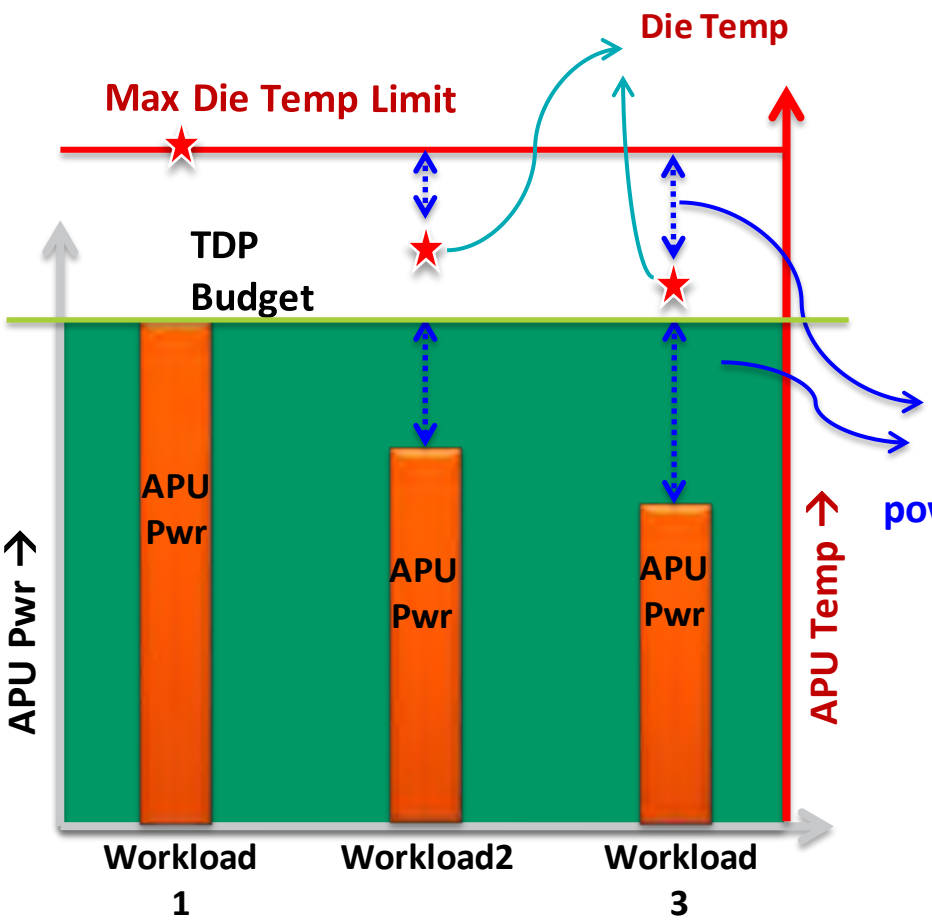

Without AMD Turbo CORE technology, power and temperature headroom may be left untapped in many workload scenarios.

Utilizes estimated and measured metrics to optimize for performance under power and thermal constraints

APU POWER MANAGEMENT

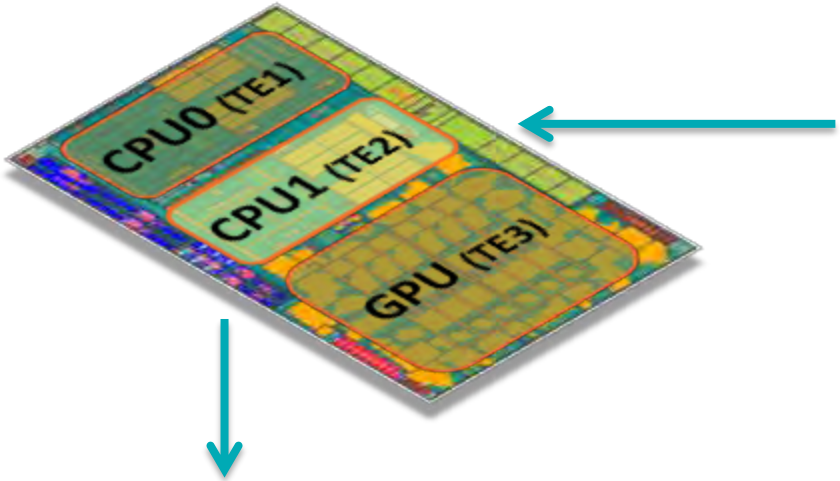

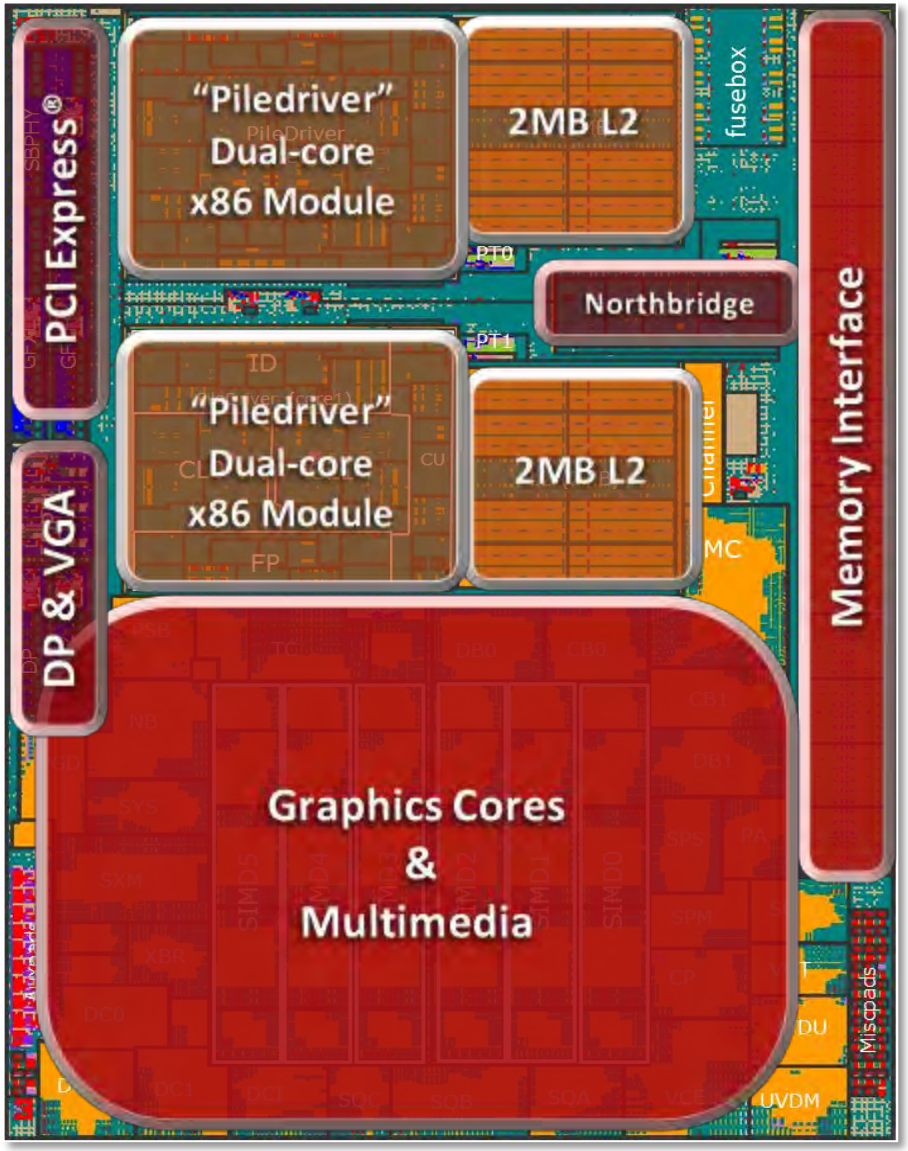

- **3 main thermal entities (TE)**
  - **TE1:** 1st x86 module + L2
  - **TE2:** 2nd x86 module + L2
  - **TE3:** Graphics + Northbridge + Multimedia
- **On each TE**
  - **Power and Temperature tracked**
  - **Frequency and Voltage controlled**
- Also account for **I/O power** influence on each of the other TEs

BUILDING BLOCKS OF AMD TURBO CORE TECHNOLOGY

POWER CALCULATION

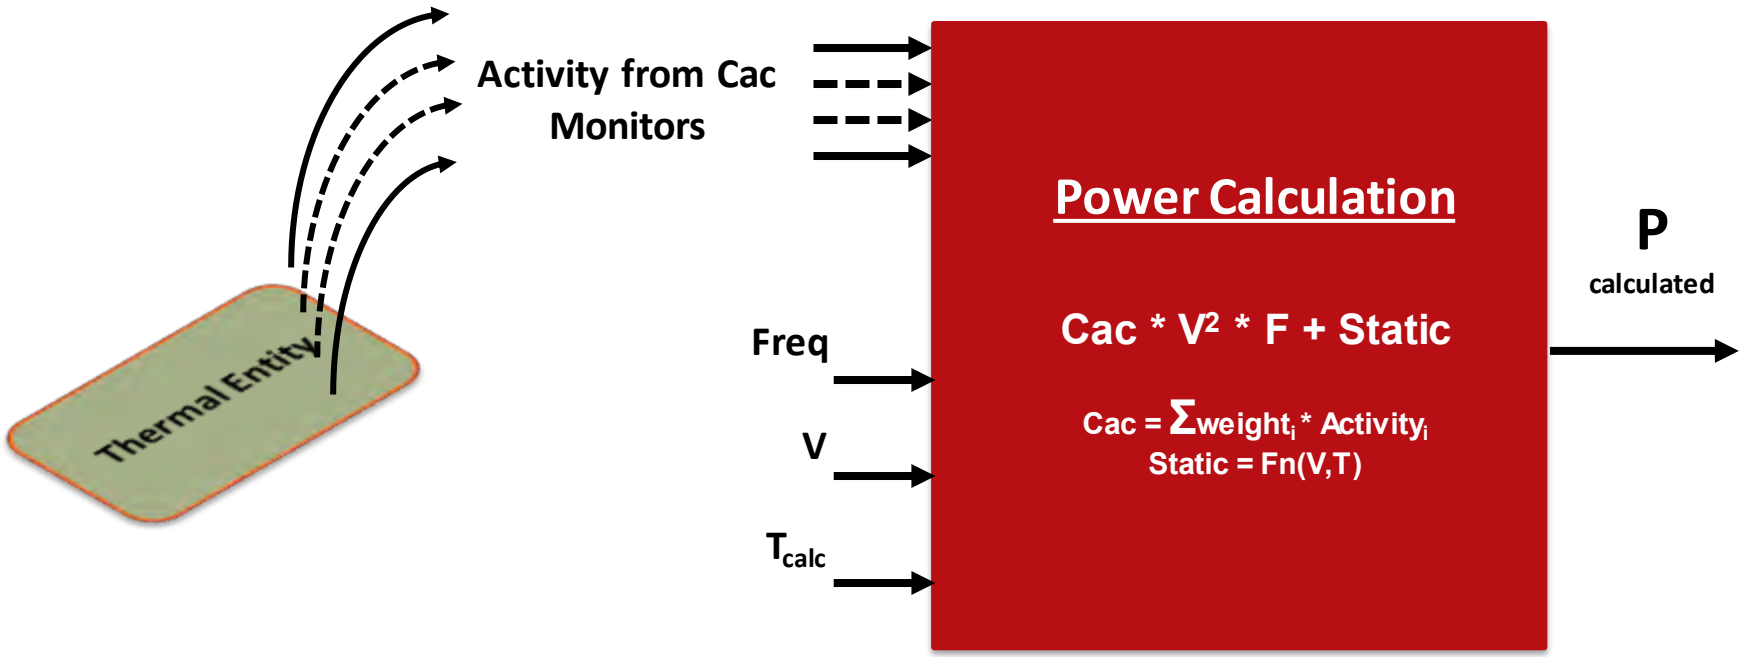

BUILDING BLOCKS OF AMD TURBO CORE TECHNOLOGY
TEMPERATURE CALCULATION

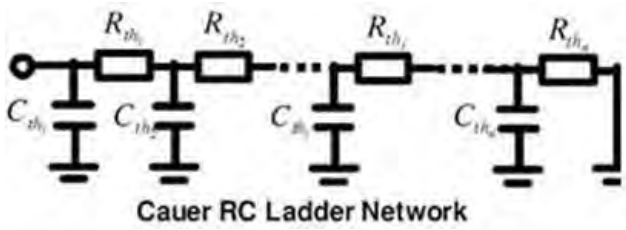

RC MODEL represents

- Cooling solution (APU heat sink, fan, etc.)
- Temperature influence of other TEs

P
calculated

Temp. Calculation

P_{calc} → **RC MODEL** → **Temp**
+
Worst-case ambient temp
+
Additional offsets
(to account for I/O power, etc.)

T
calculated

Temp.
from other TEs

Calculated temperature enables AMD to deliver robust, dependable, and repeatable performance

BUILDING BLOCKS OF AMD TURBO CORE TECHNOLOGY

P-STATE SELECTION

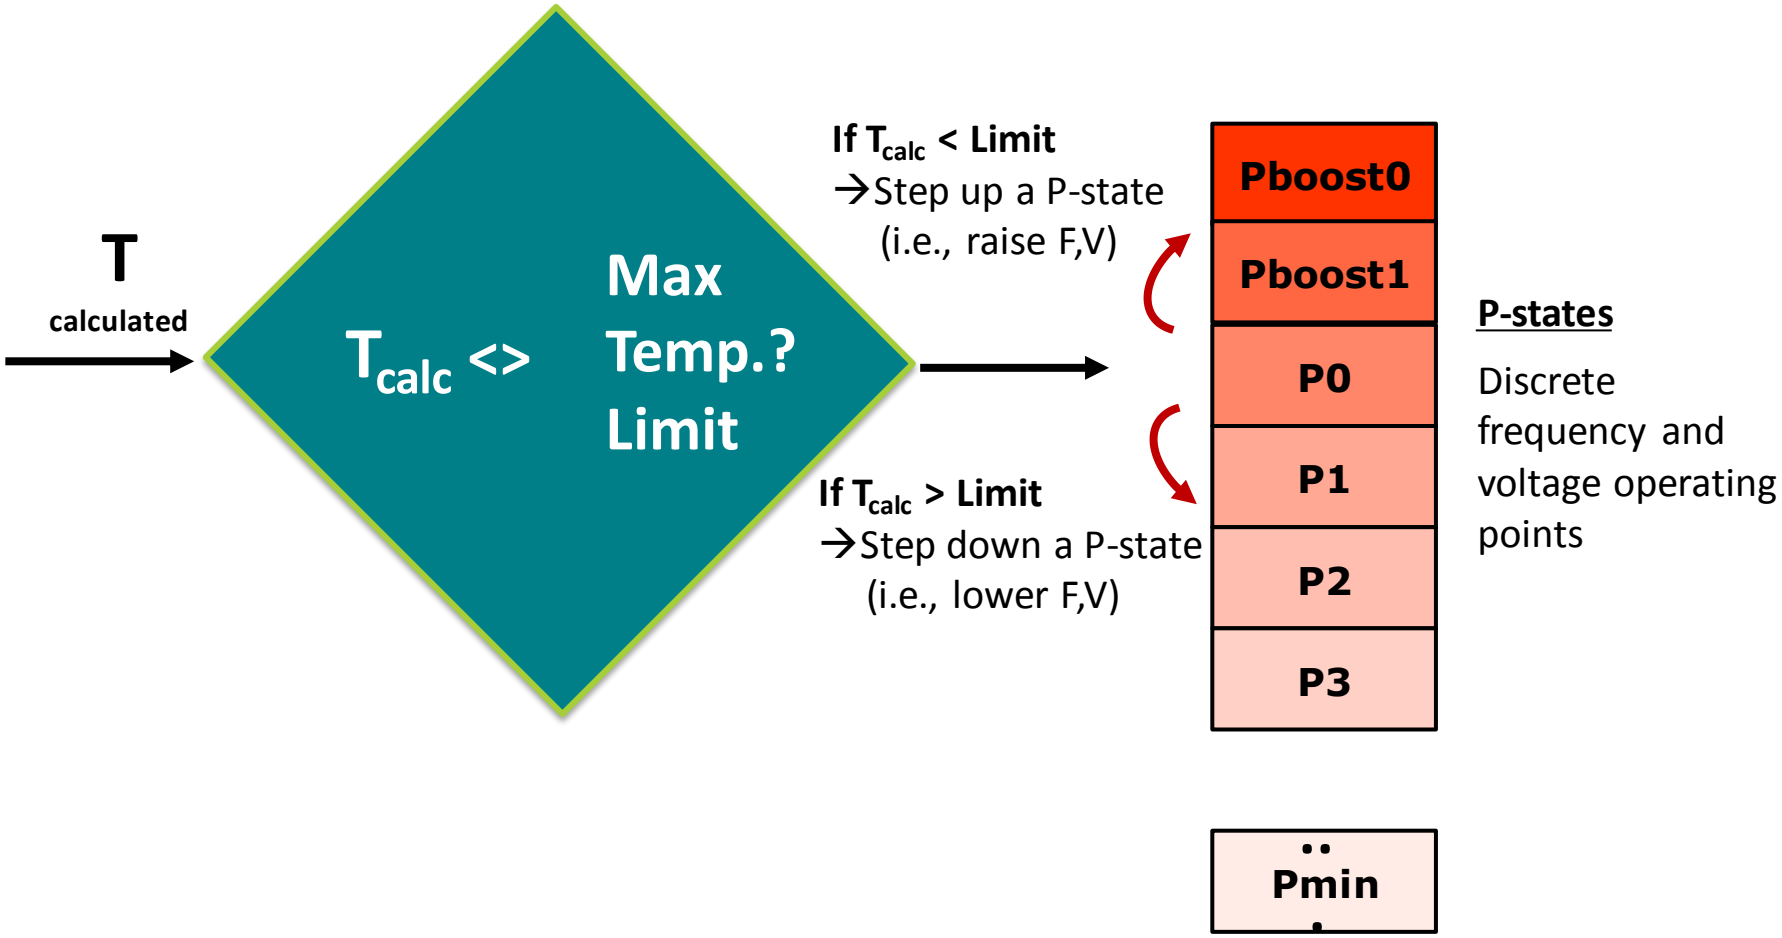

AMD TURBO CORE CONTROL LOOPS
PUTTING THE PIECES TOGETHER

AMD

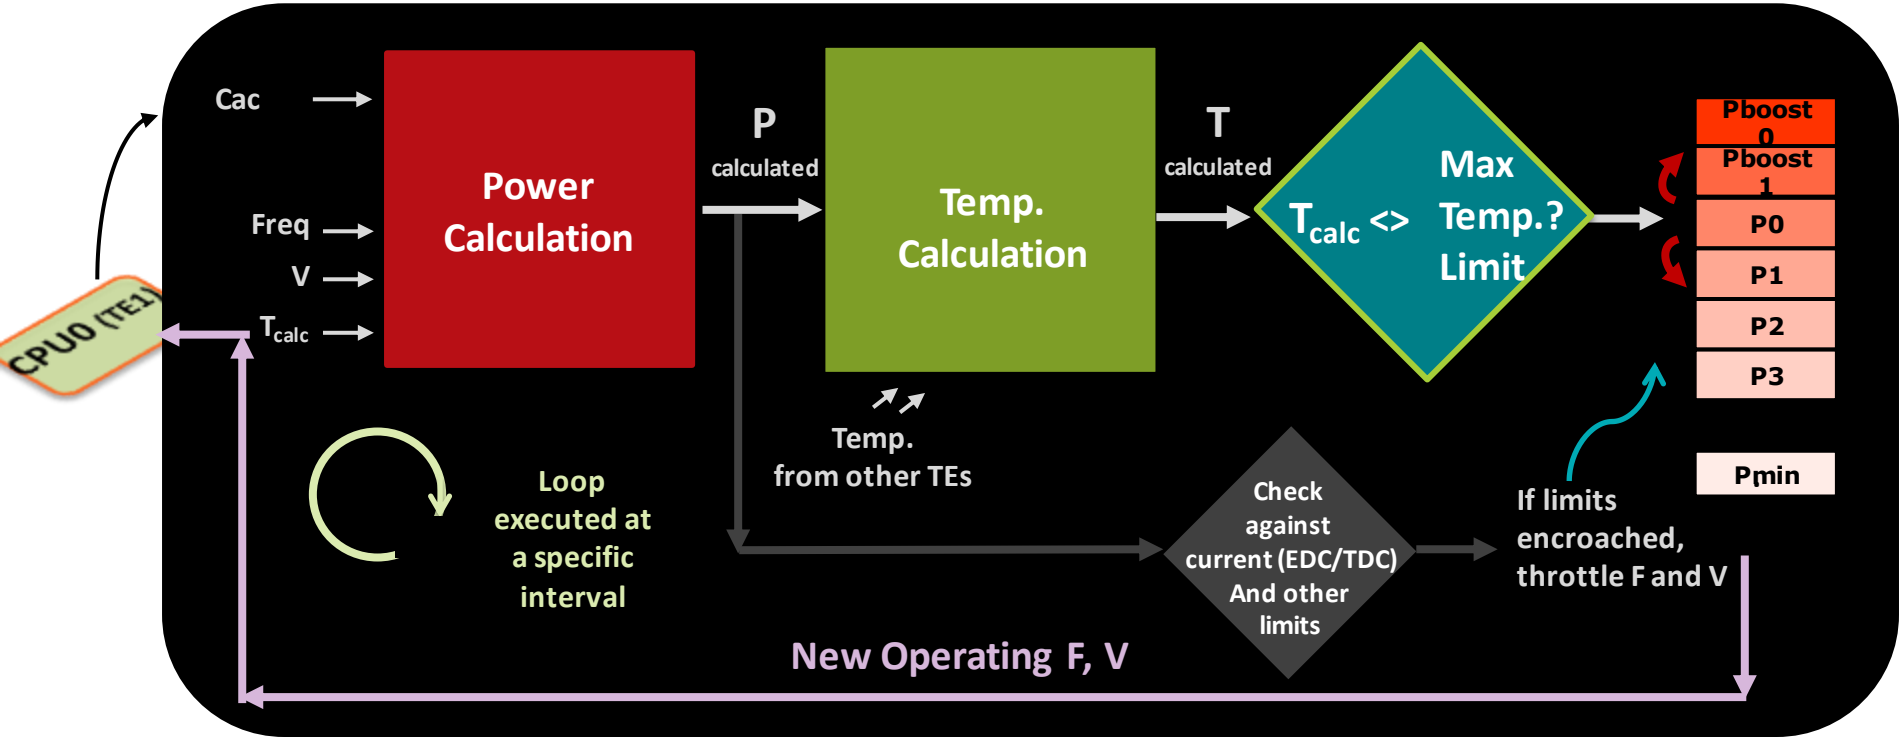

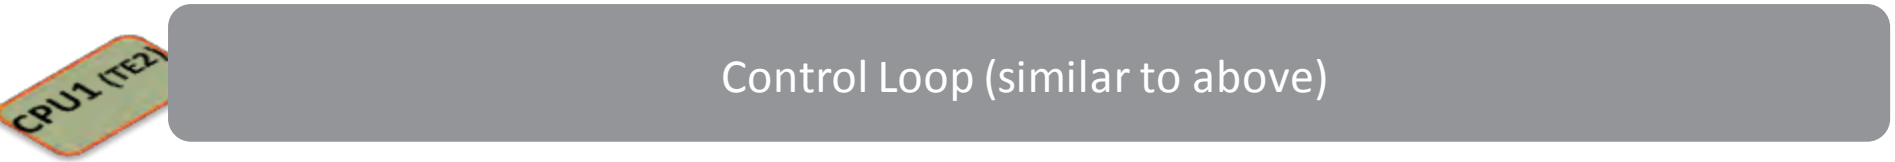

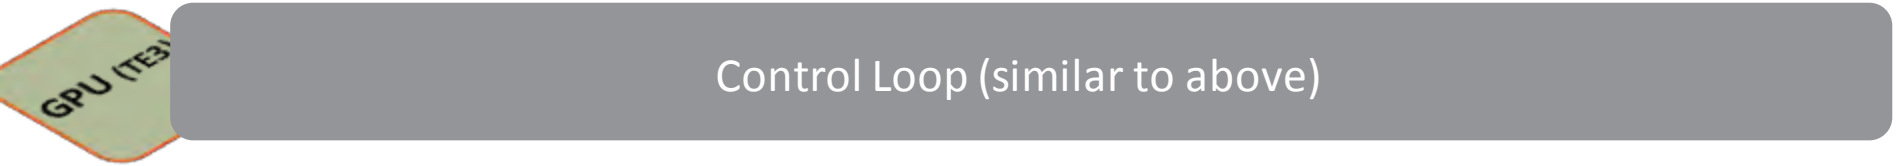

"RICHLAND" ENHANCEMENTS TO AMD TURBO CORE TECHNOLOGY

AMD

Temperature-smart AMD Turbo CORE

ADDITIONAL BOOST STATE

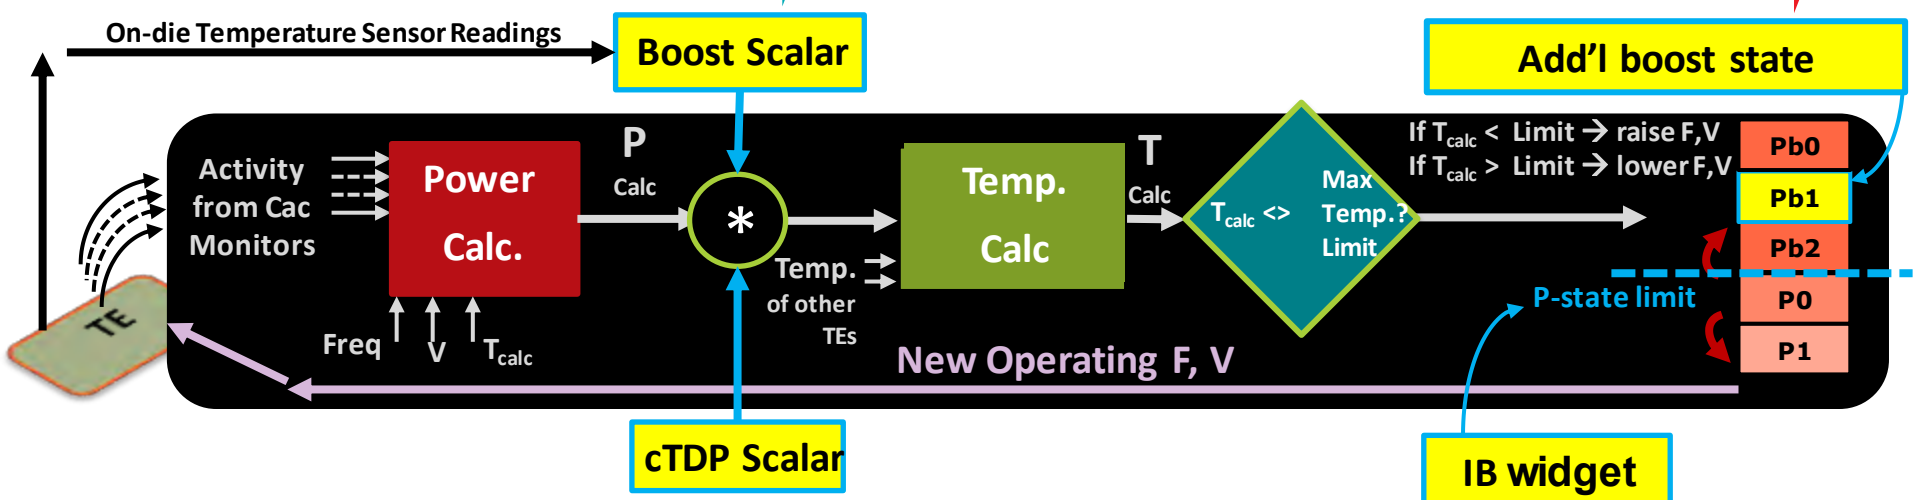

CONFIGURABLE TDP (cTDP)

Intelligent Boost (IB)

INTELLIGENT BOOST
EFFICIENT ALLOCATION OF POWER

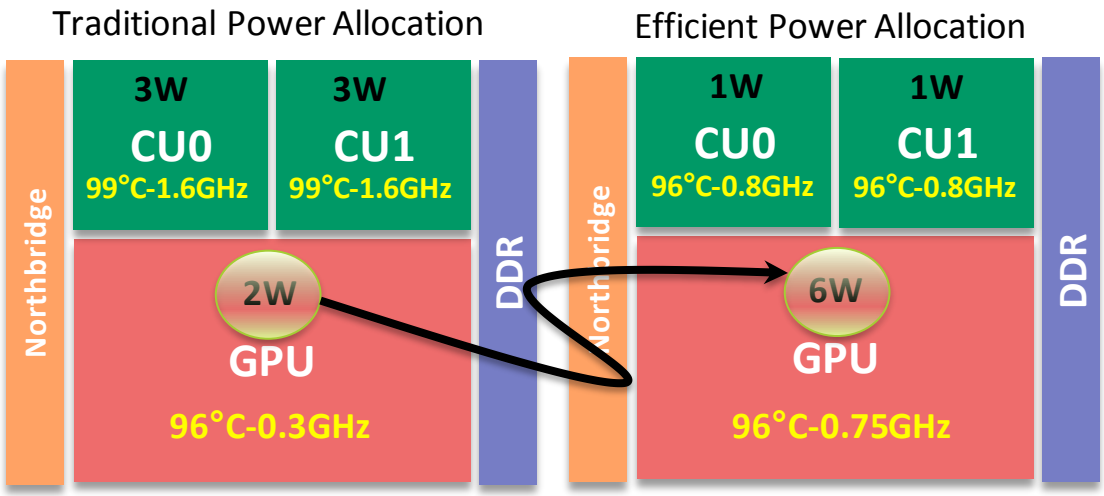

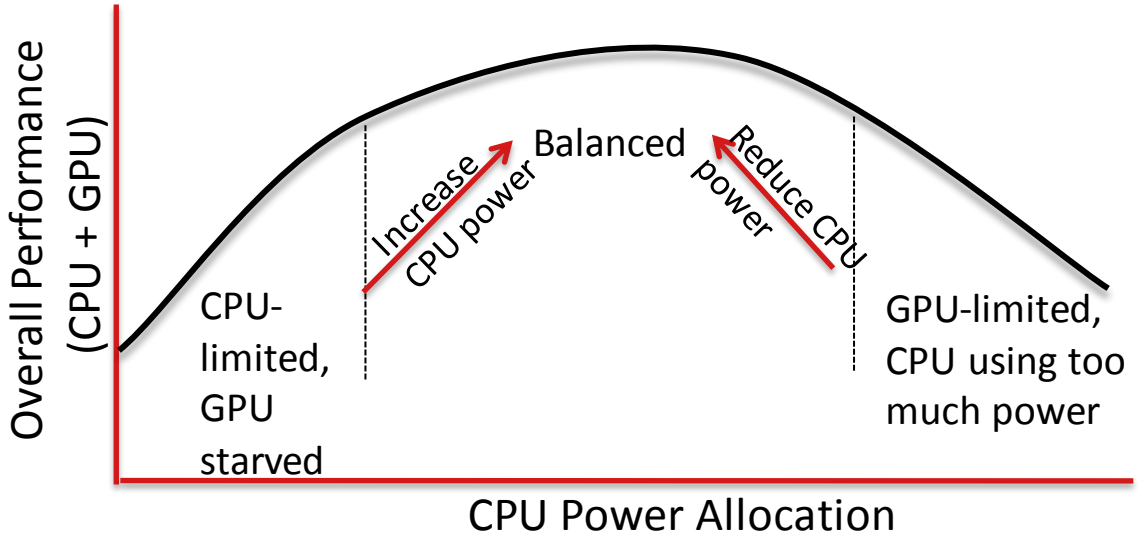

- CPU power budget at minimum level to keep GPU fully utilized
- Reduced CPU temperature designed to allow GPU to sustain higher power level
- Total system performance increases

SUMMARY

OPTIMAL UTILIZATION UNDER GIVEN PHYSICAL CONSTRAINTS

- HSA framework for effective utilization of available compute resources
  - Unified, coherent address space
  - Flexible and powerful programming interface
  - GPU as first-class compute engine

- Sophisticated power management solutions for effective operation under physical constraints
  - Extensive monitoring of internal events and activity
  - Intelligent algorithms for translating thermal headroom into performance
  - Management of power limits (cTDP)
  - Deterministic results using estimated values
  - Better margining using measured temperature values

NORTHBRIDGE AND MEMORY POWER CONTROLS

AMD

- Memory power-limiting:
  - Bandwidth capping by throttling commands
    - Rolling window of 100 clock cycles
    - Throttle by 30% to 95%
  - Controlled by
    - External feedback on thermals (DIMM temperature sensors)
    - Software register control independent of thermal management
- Memory P-states
  - 4 Northbridge P-states based on core and GPU activity levels (2 active at any time)
  - 2 memory P-states

AMD SERVER ROADMAP: JUNE 2013

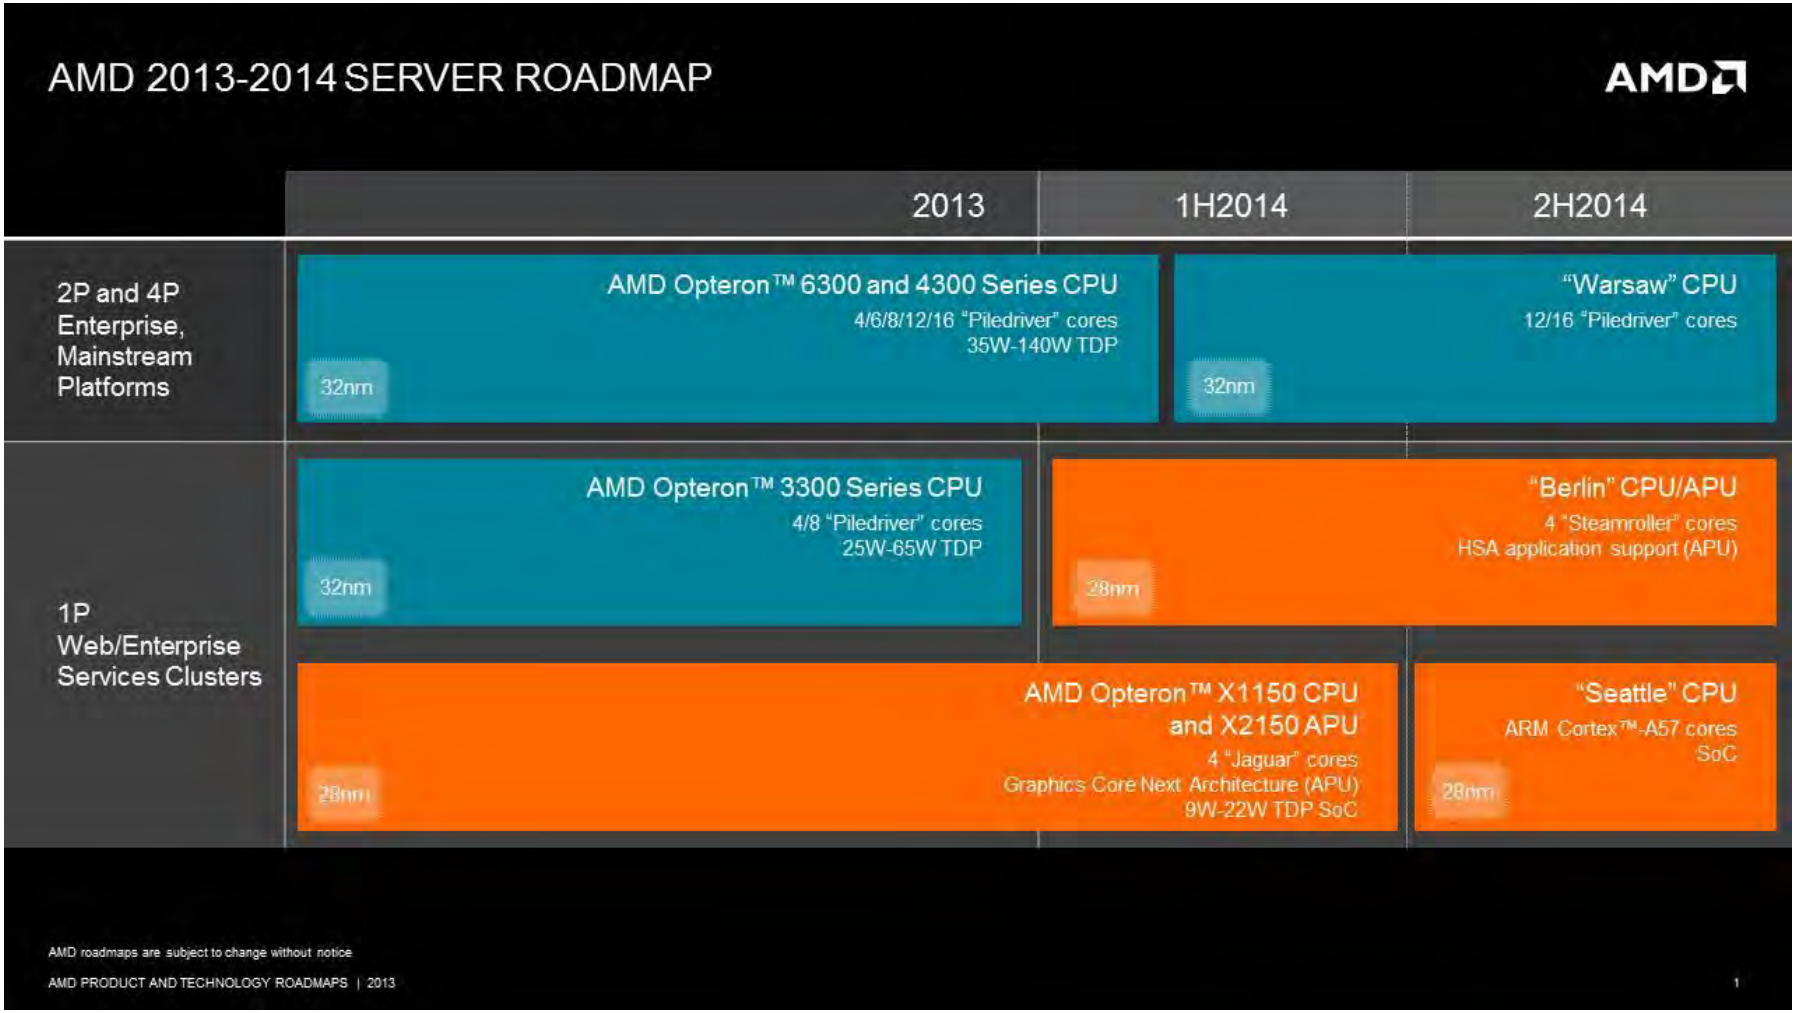

Monitoring and Management

Measurements: Nodes

(Info) A node-level measurement shall consist of a combined measurement of all components that make up a node in the architecture. For example, these components may include the CPU, memory, and the network interface. If the node contains other components such as spinning or solid state disks, they shall also be included in this combined measurement. The utility of the node-level measurement is to facilitate measurement of the power or energy profile of a single application. The *node* may be part of the network or storage equipment, such as network switches, disk shelves, and disk controllers.

(important) The ability to measure the current and voltage of any and all nodes shall be provided.

The current and voltage measurements shall provide a readout capability of:

(mandatory) ≥ 1 per second
(important) ≥50 per second
(enhancing) ≥ 250 per second

(mandatory) The current and voltage data must be real electrical measurements, not based on heuristic models.

VENDOR FORUM SLIDE

APML: ADVANCED POWER MANAGEMENT LINK

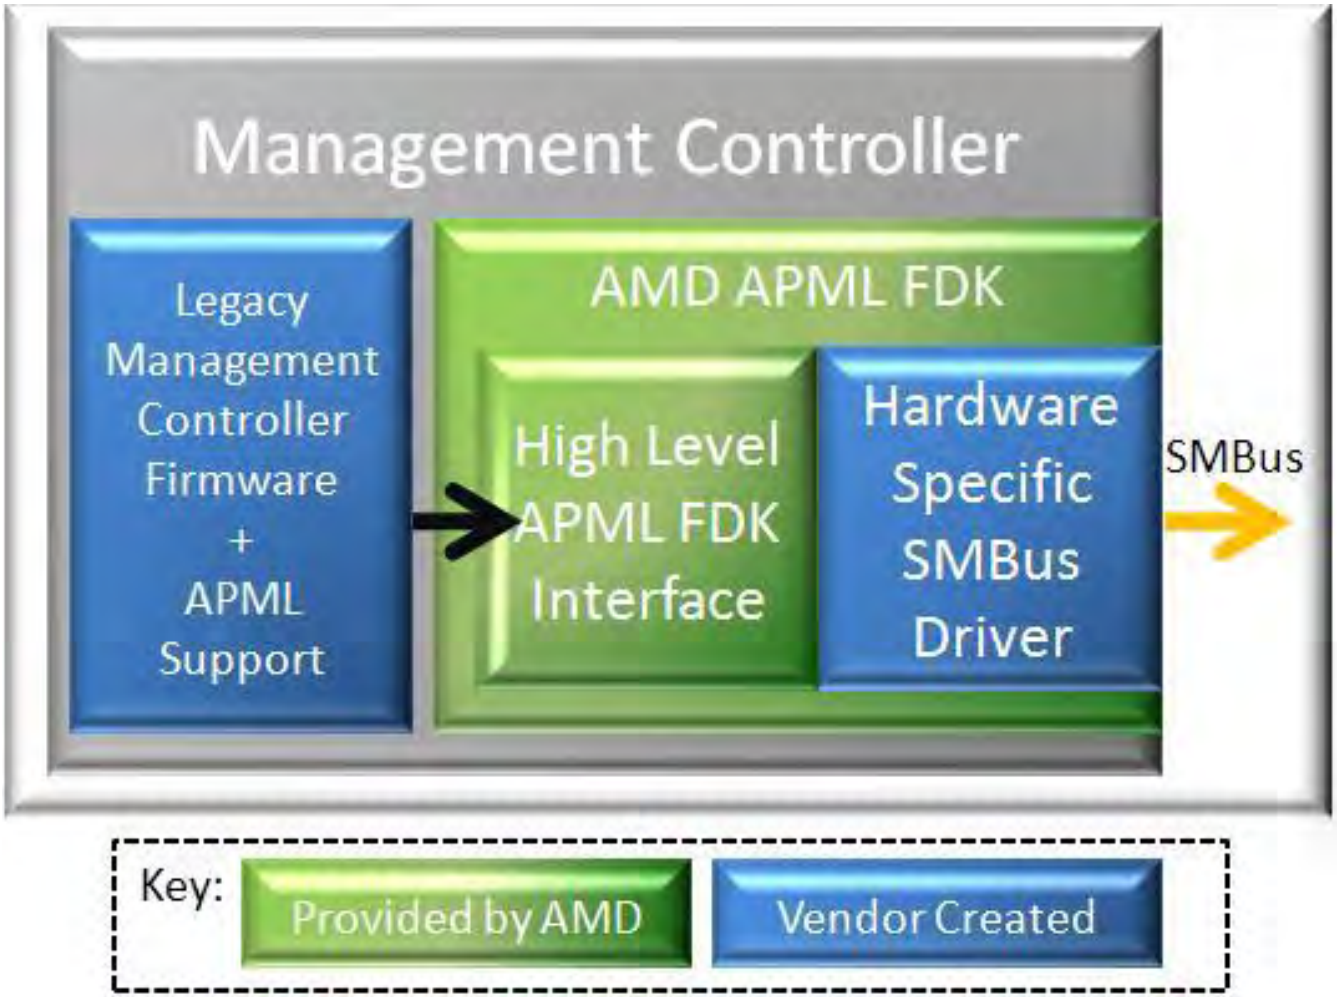

APML FEATURES

- Examples:
  - Processor power-capping
  - Processor asset identification (including CPUID)
  - Machine check register access with optional alerts to management subsystems
- Specifics
  - 100 kHz standard, and 400 kHz fast mode
  - SB-TSI: Sideband temperature sensor interface
    - Access internal temperature sensor and specify temperature thresholds
  - SB-RMI: Sideband remote management interface
    - Monitor the current P-state
    - Set and read current maximum P-state

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

© 2013 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo and combinations thereof are trademarks of Advanced Micro Devices, Inc. in the United States and/or other jurisdictions. Other names are for informational purposes only and may be trademarks of their respective owners.

Backup