Architectures for HPC Workloads – Trends in Microprocessor development

Leif Nordlund

HPC Business Development EMEA

August 2010

X86 purchase criteria Changed Over Time

Over time, the criteria has changed from:

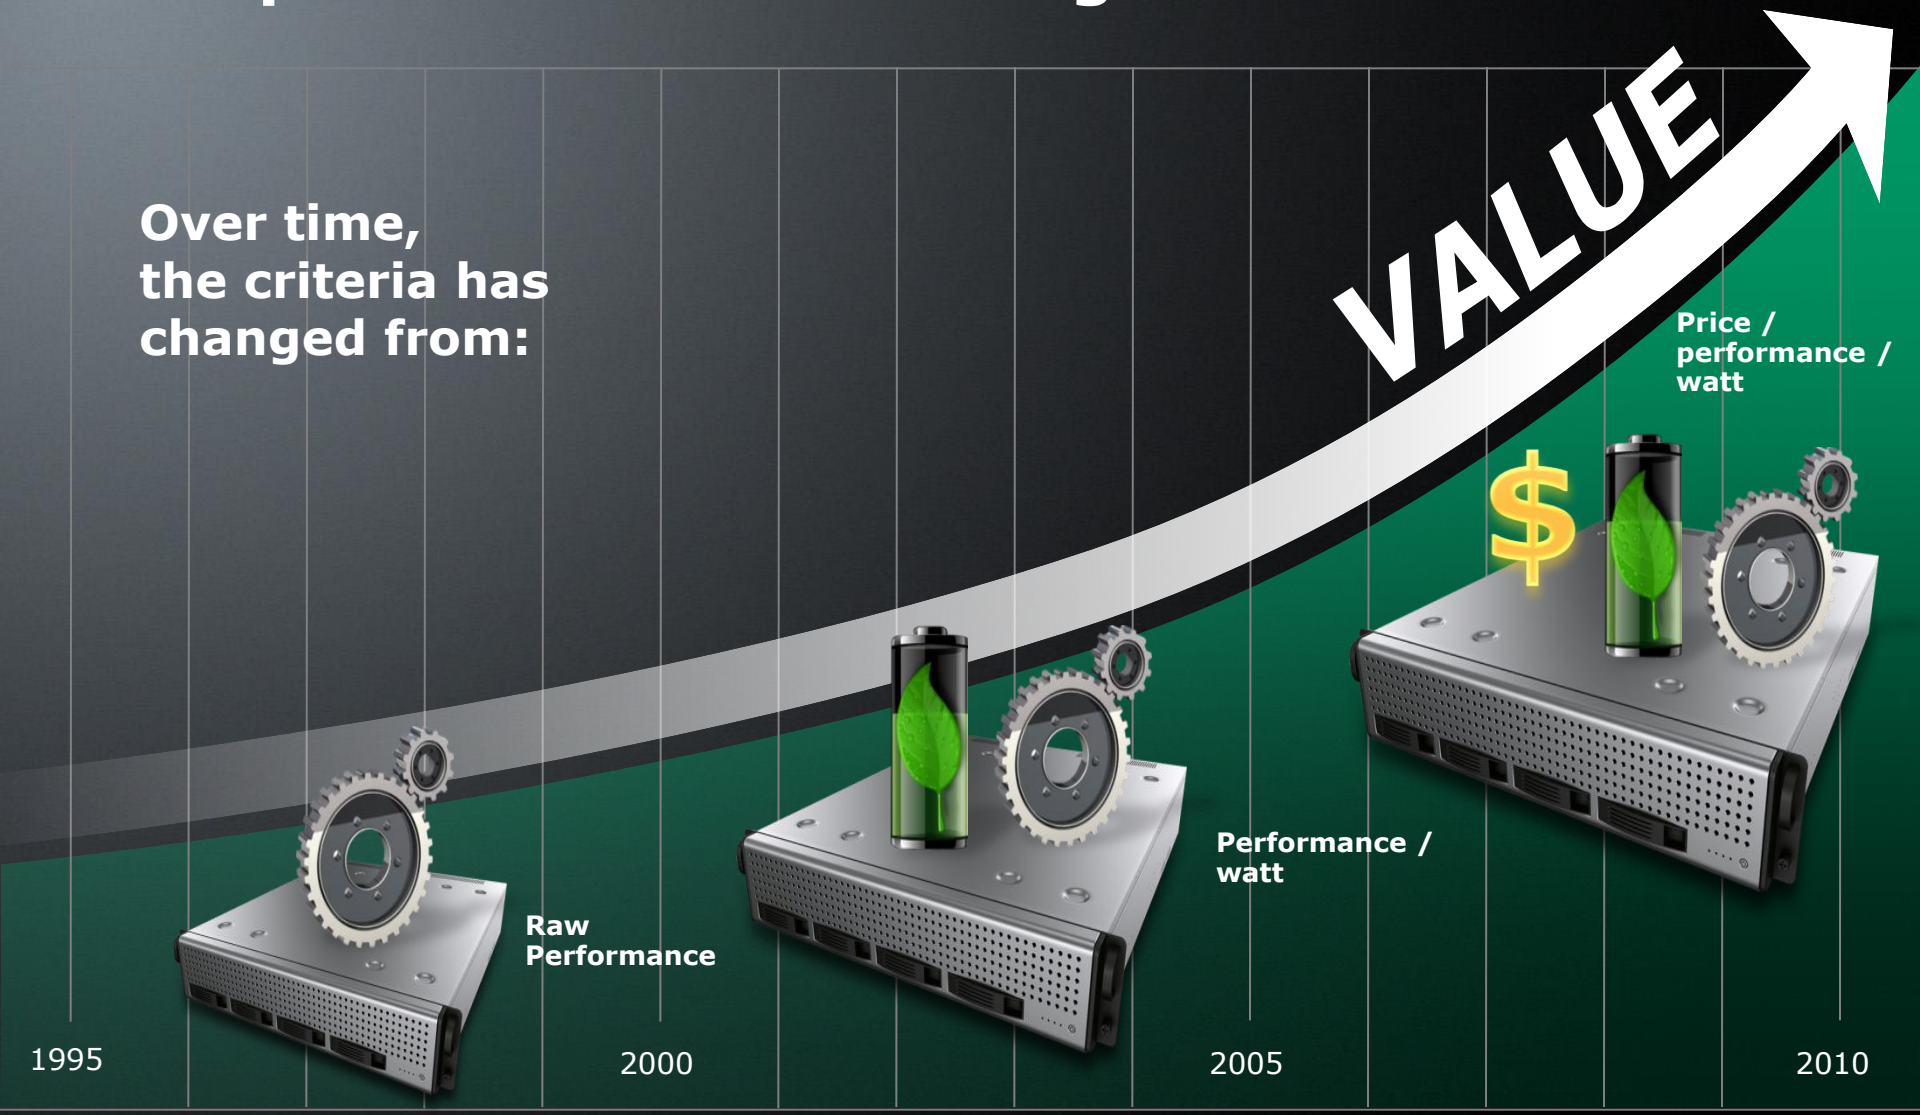

The road to HPC/Exascale using volume products

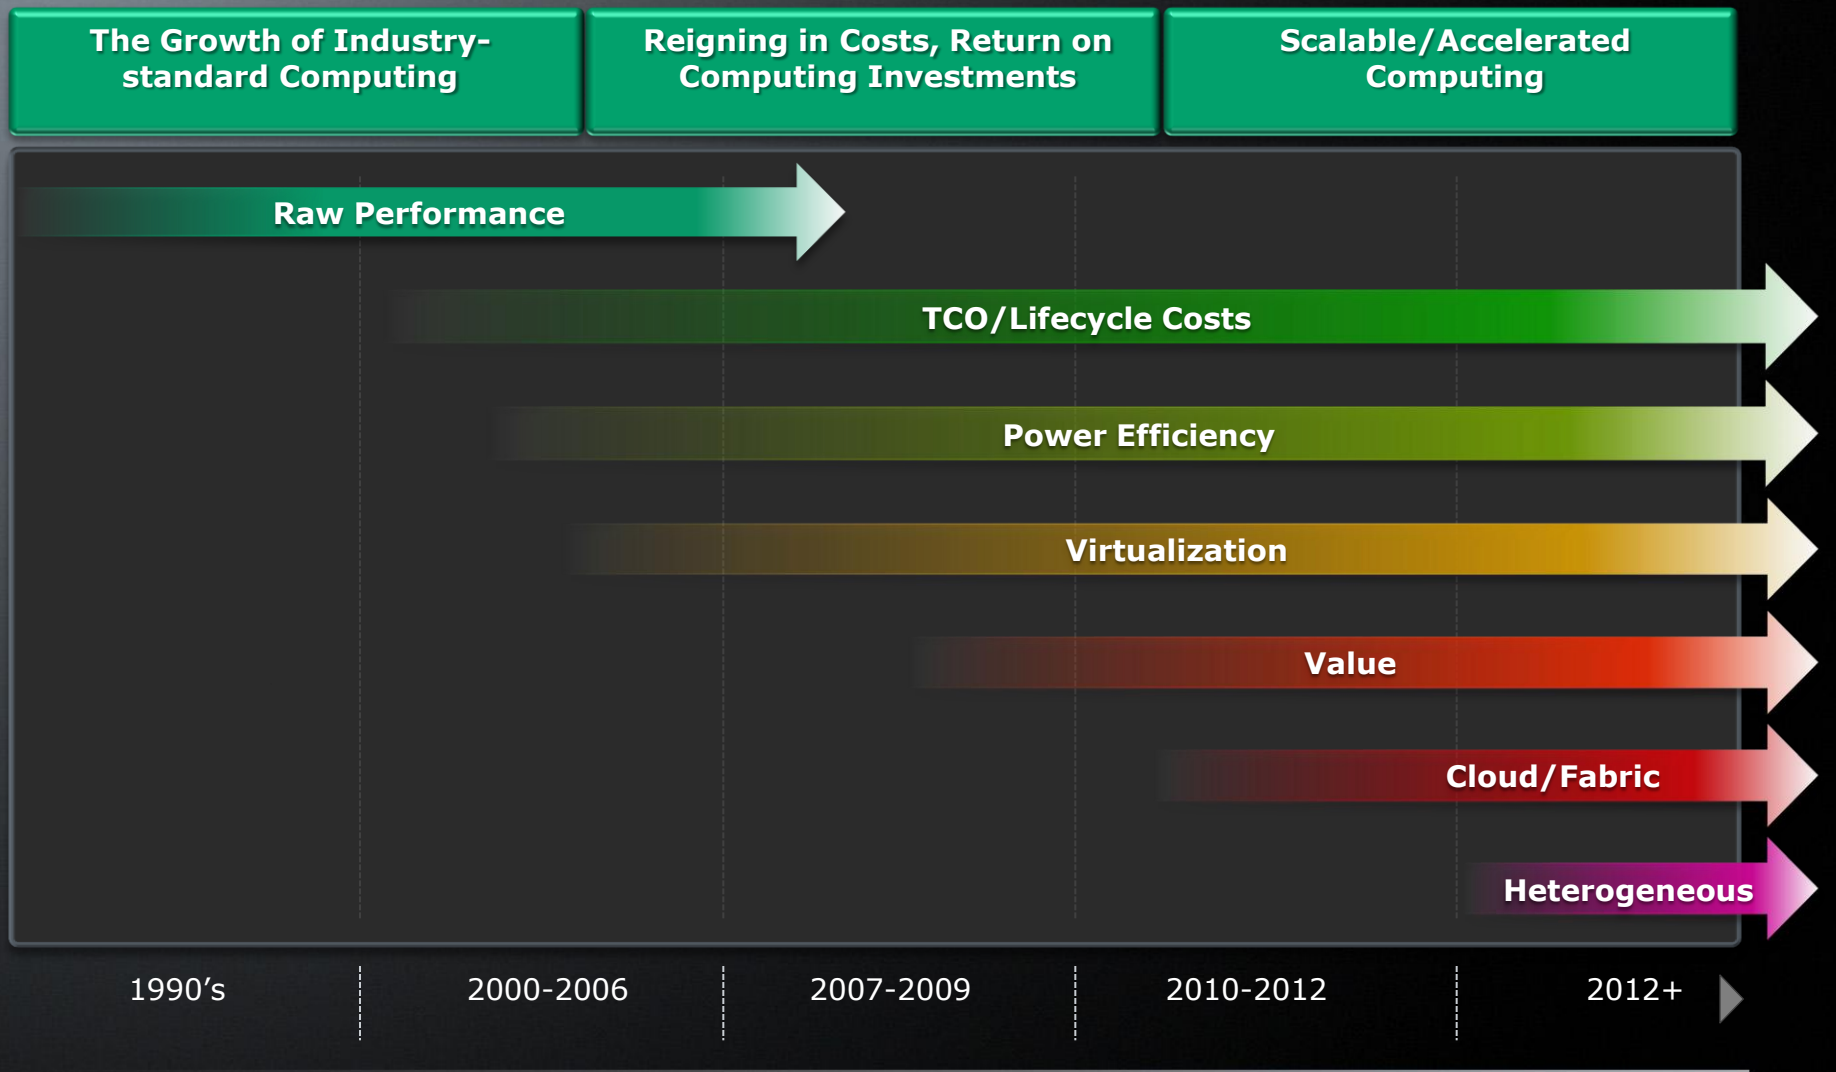

Current Microprocessors / Volume products

x86 CPUs and Graphics GPUs for PCs (and Servers)

Multi-core x86 Processors

- Performance/Watt using up to 12 cores
- X86 instruction set (+extensions)
- Enhanced Efficiency and instructions

GPUs for Graphics

- 3D Accelerators For Visualization
- See More and Do More with Your Data

GPU Computing (GPGPU)

- GPU Optimized For Computation
- Massive Data-parallel Processing
- High FLOP Performance Per Watt

AMD FireStream™ 9350
Maximum compute density GPGPU (4GPU in 1U)

Breakthrough performance in a single-slot solution

2.0+ TFLOPS single precision floating point

400 GFLOPS double precision floating point

2GB GDDR5 memory

Single slot with passive heat sink

<= 150W peak TDP
(single 6-pin AUX connector)

Delivery in Q3 2010

MSRP $799

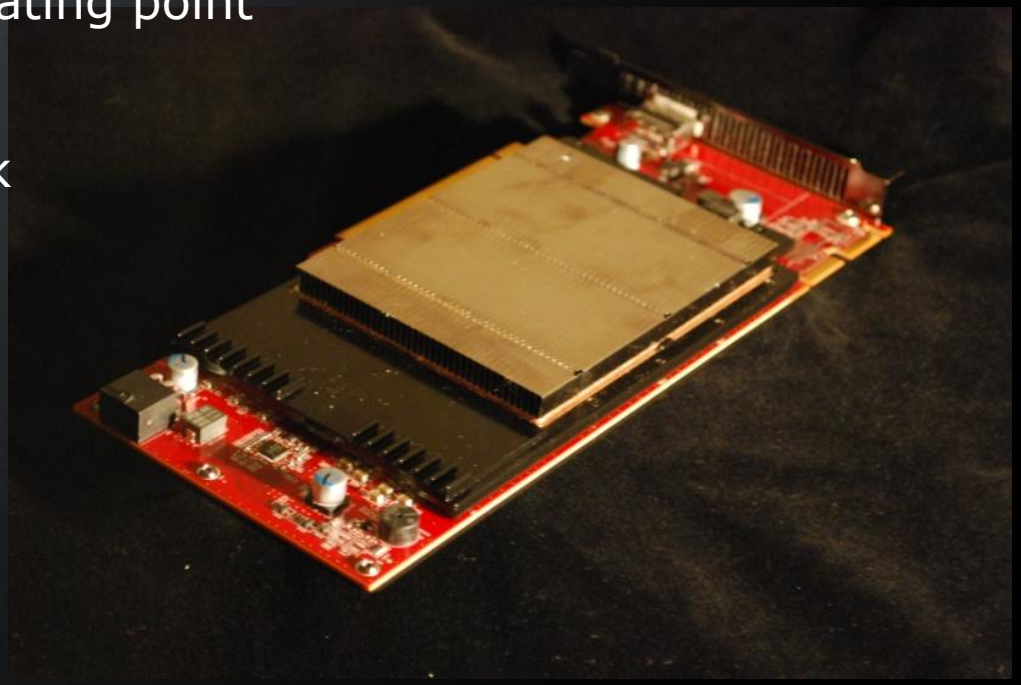

Enhancing the x86 core/watt and Instruction set

The AMD "Bulldozer" core has shared and dedicated components

The shared components:

–Help reduce power consumption

–Help reduce die space

The dedicated components:

–Help increase performance and scalability

Bulldozer dynamically switches between shared and dedicated components to maximize performance per watt

One of the building blocks of the next generation of FUSION processor designs

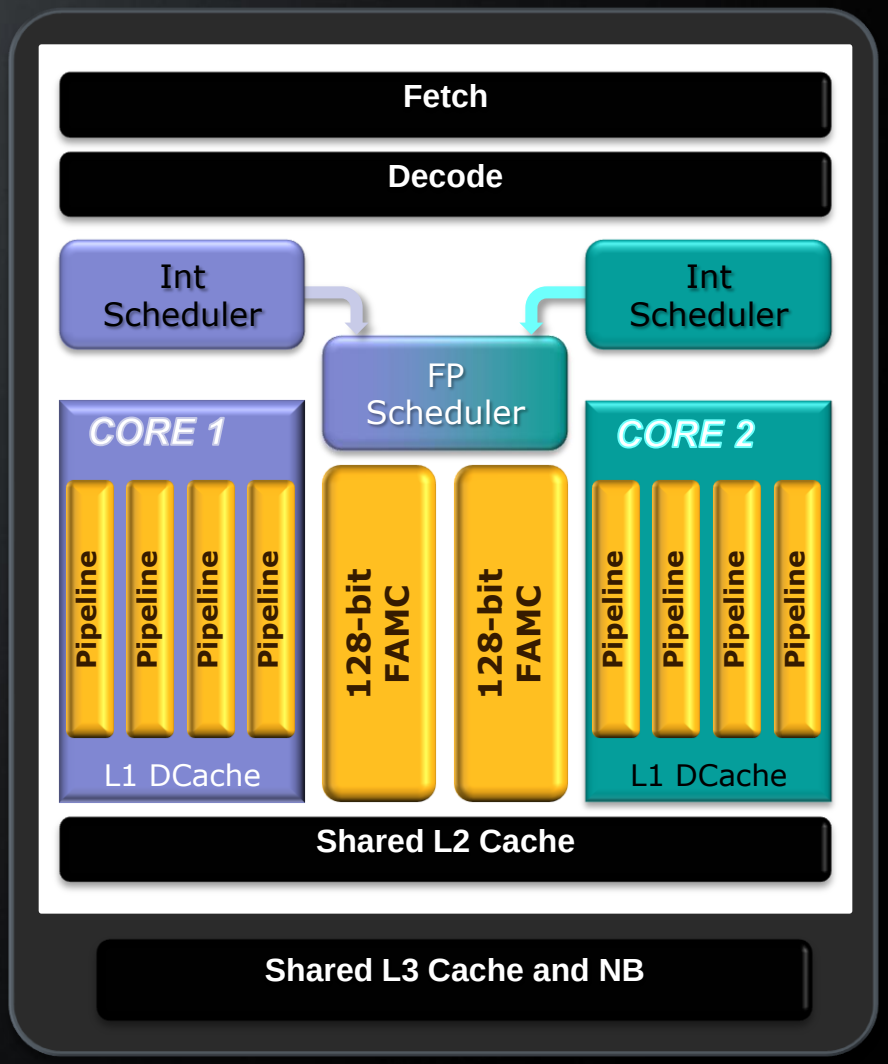

Designed for Scalability and Performance

"Bulldozer" module

Two cores in a single unit that enables two simultaneous threads, the building blocks of a "Bulldozer" die

Parallel Threads

The ability to execute two threads on two discrete, unshared cores without compromising or increasing bottlenecks

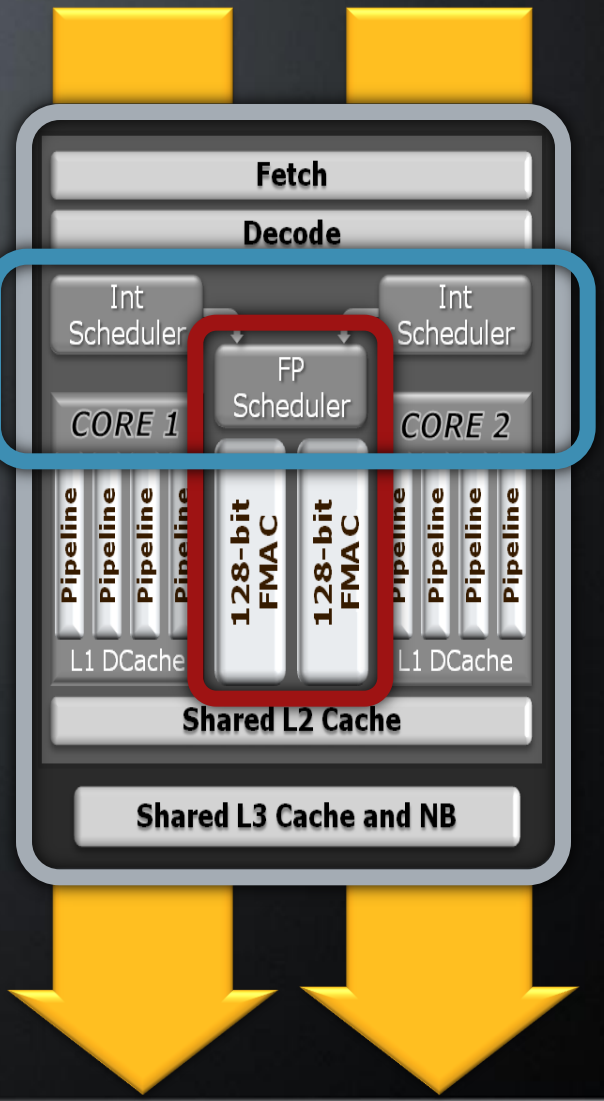

Flex FP

A flexible floating point unit that can be dedicated OR shared between the two cores per cycle

Dedicated Scheduler

Independent integer schedulers and an FP scheduler help improve scalability by efficient execution

ISA For First-generation Bulldozer

Specification: http://support.amd.com/us/Processor_TechDocs/43479.pdf

More background: http://forums.amd.com/devblog/ - "Striking A Balance"

Achieves parity with Intel in instruction set extensions

- SSSE3
- SSE4.1
- SSE4.2
- AES (acceleration for AES encryption standard)
- PCLMULQDQ (carryless multiply -- crypto, CRC acceleration)
- AVX "Advanced Vector Extensions" – with 256 bit SIMD Vector registers

Plus the following AMD extensions

- XOP - e**X**tended **OP**eration: multimedia and vectorization extensions beyond AVX
- FMA4 - four-operand Fused Multiply Add
- LWP - Lightweight Profiling - dynamic performance optimization support

AMD "APU" FUSION Integrated Chip – coming 2011

Combination of CPU and programmable GPU architectures for high-performance heterogeneous compute capability

- High-speed bus architecture
- Shared, low-latency memory model
- Single die design

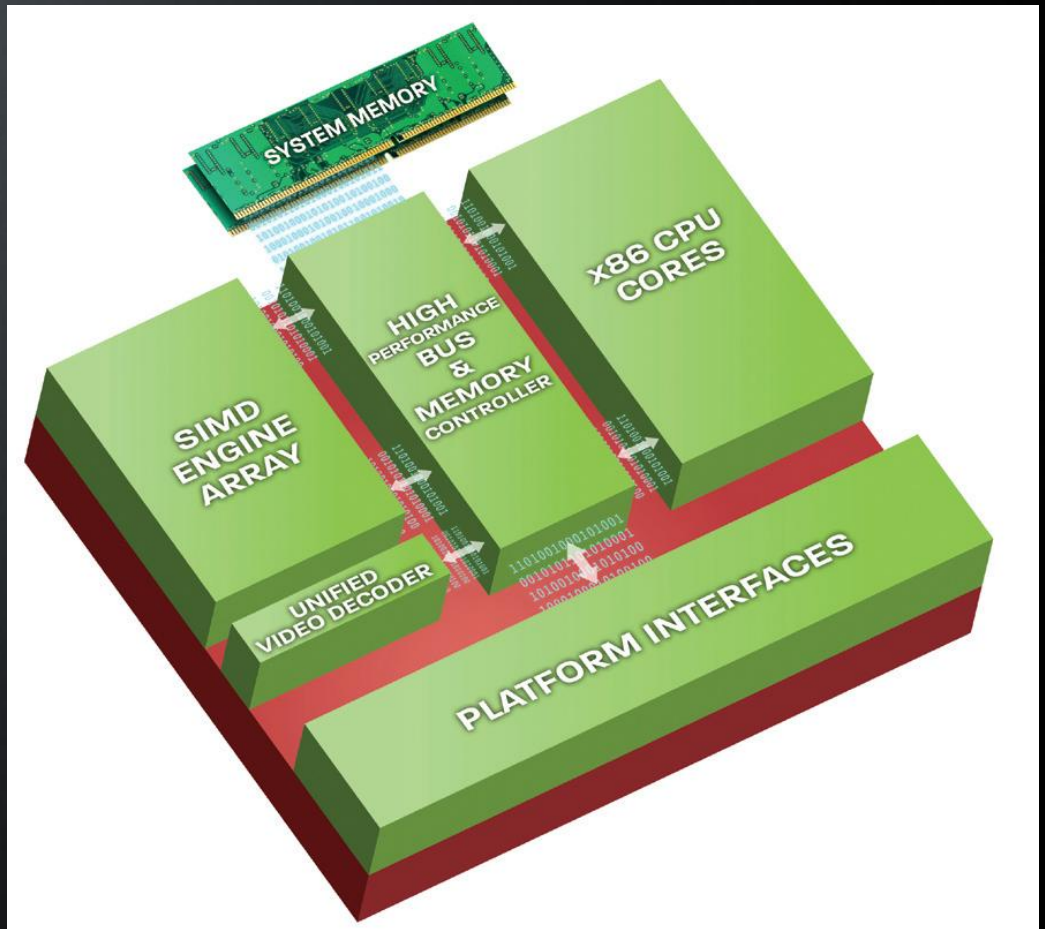

The Road to Server Fusion

AMD's multi-year enablement strategy brings GPGPU together with AMD Fusion architecture to help accelerate technical computing

Integration Steps:

1. Application level tools enable GPGPU computing
2. Drivers deliver operating system support
3. PCI integration in CPU delivers a direct communication link
4. Unified communications; DMA memory and cache access w/accelerators
5. Integration of CPU, GPU and communications with shared memory model

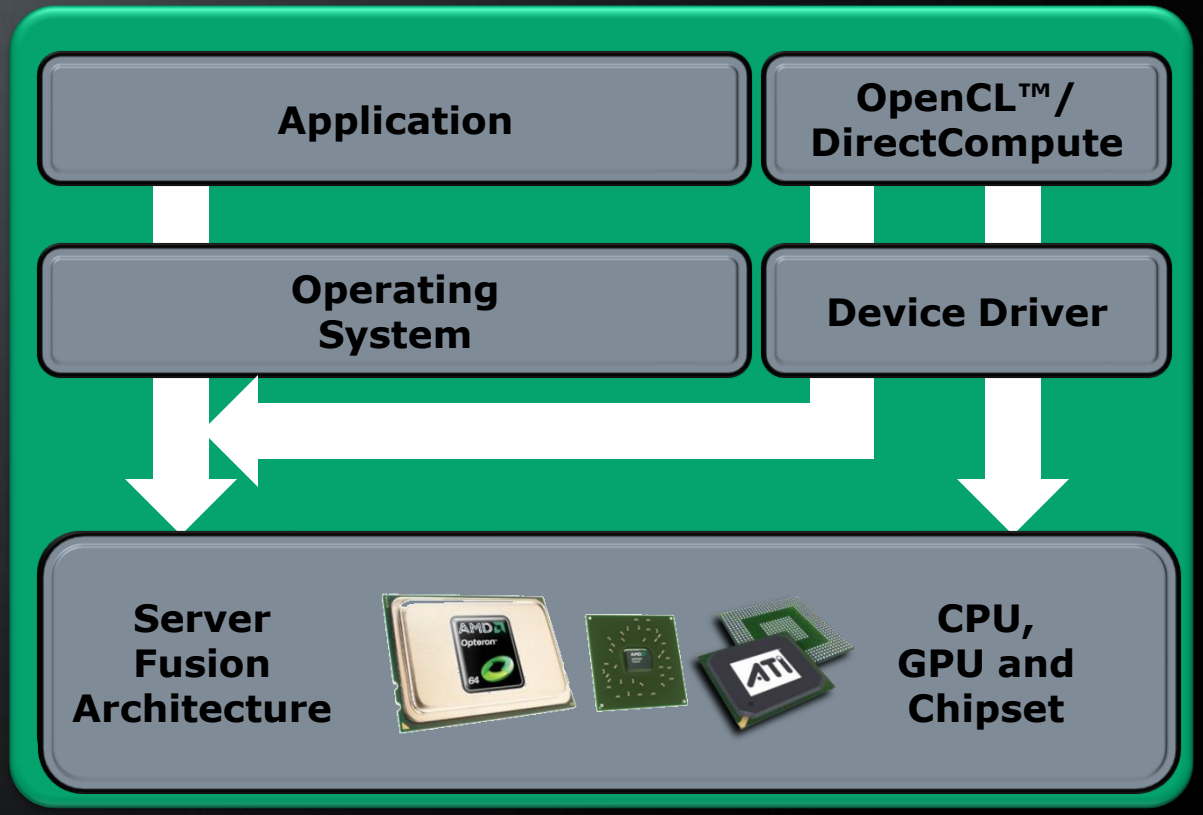

Heterogeneous Computing

Industry Standards Drive Adoption

2008
Proprietary Research

Brook+ | CUDA | C

ATI GPUs | Nvidia GPUs | x86

- Developer interacts with proprietary GPGPU APIs and HW
- Separate API for CPUs
- Applications hand-tuned for target architectures

2010
Industry Standards

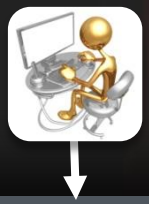

OpenCL™*

GPUs | x86

- OpenCL scales on both GPUs and CPUs – developers can balance workloads using single API
- Streamlined cross-platform development

2012
Fusion

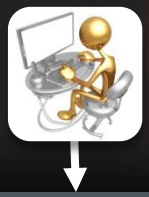

IDEs

Libraries, Tools

OpenCL™

GPUs | x86 | APUs

- Fusion architecture and ecosystem support dynamic load balancing
- Broad application adoption

Hardware abstraction

*OpenCL and the OpenCL logo are trademarks of Apple Inc. used by permission by Khronos.

ONLY AMD!

CPU

GPU

Only

AMD

ATI STREAM TECHNOLOGY

OpenCL

KHRONOS GROUP

fusion

Disclaimer and Attribution

CAUTIONARY STATEMENT
This release contains forward-looking statements concerning revenue and other expectations which are made pursuant to the safe harbor provisions of the Private Securities Litigation Reform Act of 1995. Forward-looking statements are commonly identified by words such as "would," "may," "expects," "believes," "plans," "intends," "projects," and other terms with similar meaning. Investors are cautioned that the forward-looking statements in this release are based on current beliefs, assumptions and expectations, speak only as of the date of this release and involve risks and uncertainties that could cause actual results to differ materially from current expectations.