AMD
Smarter Choice

CONFIDENTIAL

Future of Supercomputing: The Computational Element

Tommy Toles
512-327-5389
tommy.toles@amd.com

Agenda

AMD
Smarter Choice

- **HPC – The New World Order**
- **Whole > Sum-of-Parts?**
- **Key Challenges**
- **A Look Ahead**

All Worldwide Servers Compared To HPC Revenue, Units & Processors

AMD
Smarter Choice

All Servers Worldwide						
	2003	2004	2005	2006	2003 to 2006 CAGR	2005 to 2006 CAGR
Total Factory Revenue ($B)	$46,149	$49,146	$51,268	$52,251	4.2%	1.9%
Units Shipped (same as nodes)	5,278,222	6,307,484	7,050,099	7,472,649	12.3%	6.0%
Processor Dies Shipped	8,662,823	10,134,624	11,712,766	12,779,159	13.8%	9.1%
Source: IDC 2007						
HPC Technical Servers Worldwide						
	2003	2004	2005	2006	2003 to 2006 CAGR	2005 to 2006 CAGR
HPC Server Revenue ($B)	$5,698	$7,393	$9,208	$10,030	20.7%	8.9%
Adjusted Revenues (To match ente	$5,128	$6,654	$8,287	$9,027	20.7%	8.9%
Node Units Shipped	411,327	734,510	1,215,735	1,419,221	51.1%	16.7%
Processor Elements Shipped	1,002,905	1,657,827	2,681,079	3,351,843	49.5%	25.0%
Source: IDC 2007						
HPC As A Ratio Of All Servers						
	2003	2004	2005	2006		
Revenue ($B)	12.3%	15.0%	18.0%	19.2%		
Adjusted Revenues (Apples-to-appl	11.1%	13.5%	16.2%	17.3%		
Units Shipped (Nodes)	7.8%	11.6%	17.2%	19.0%		
Processors Shipped	11.6%	16.4%	22.6%	26.1%		
Source: IDC 2007						

Total HPC Revenue by Processor Type

Source: IDC

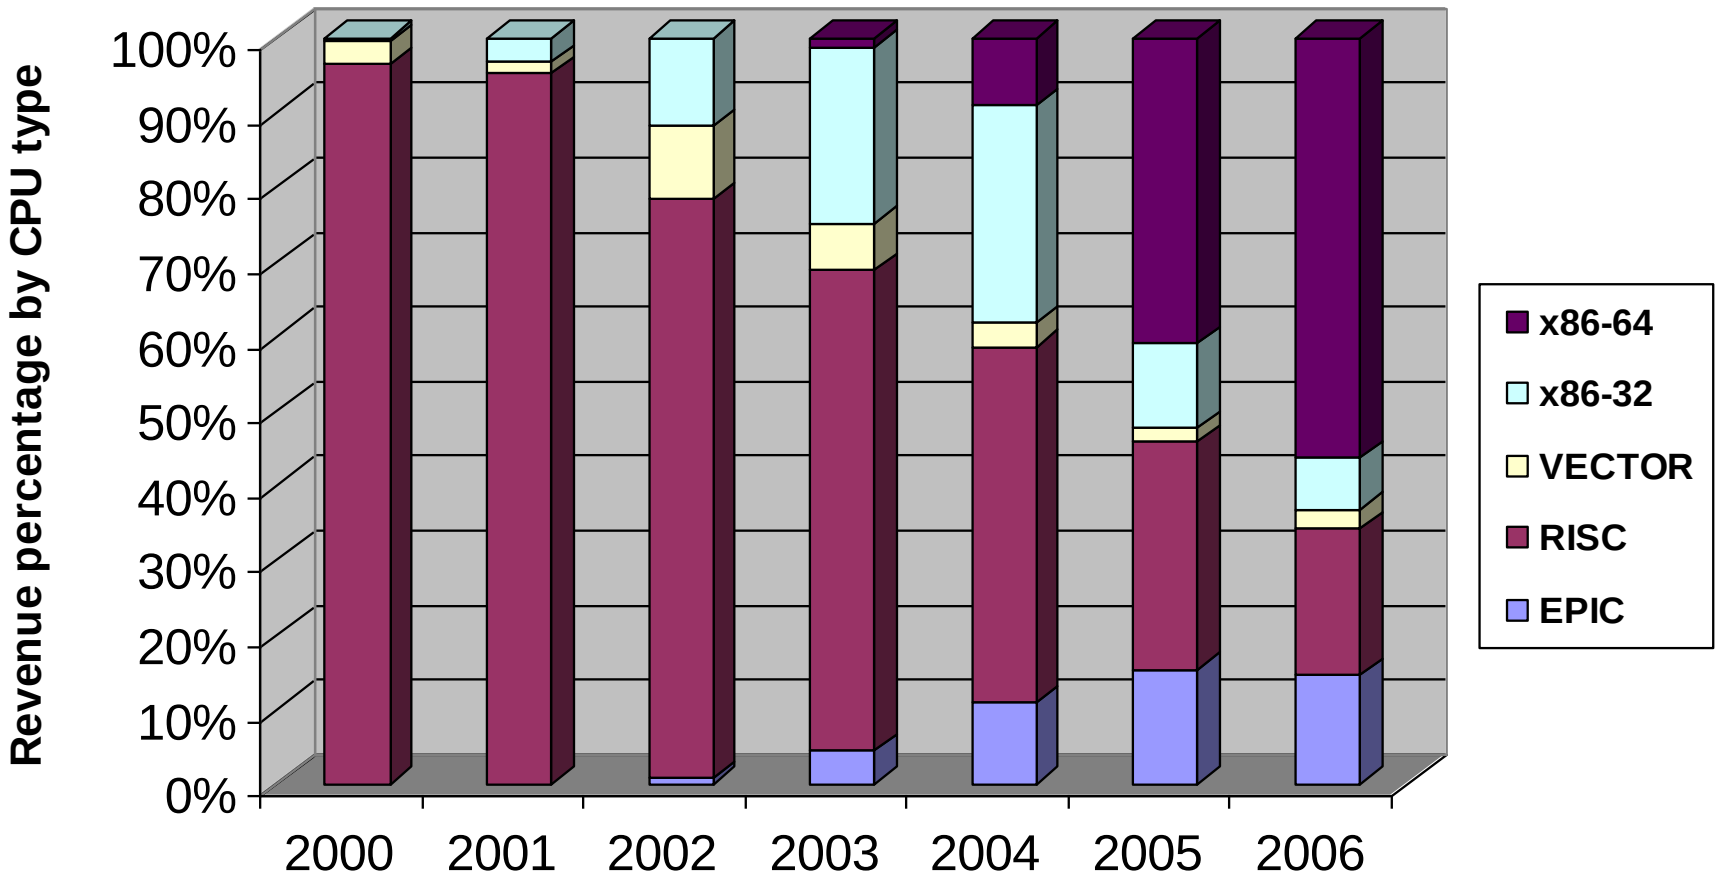

X86 bases system took 63% share in 2006

Total HPC Revenue by OS

Source: IDC

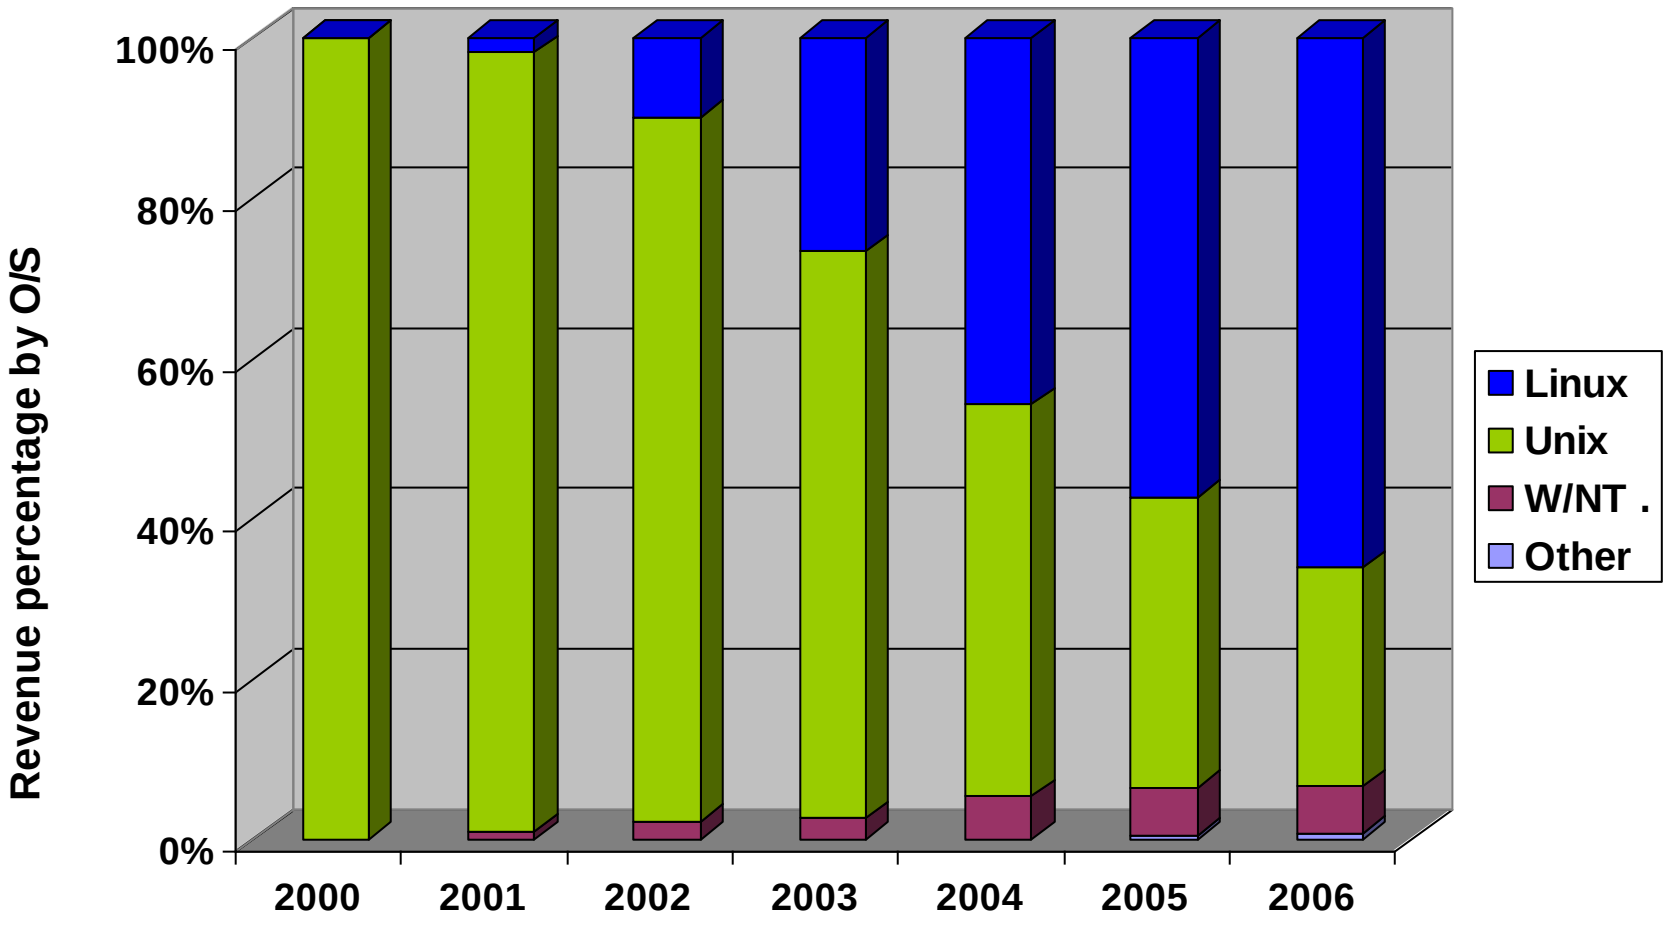

Linux systems accounted for 66% of the total revenue in 2006

Historical Performance Metrics

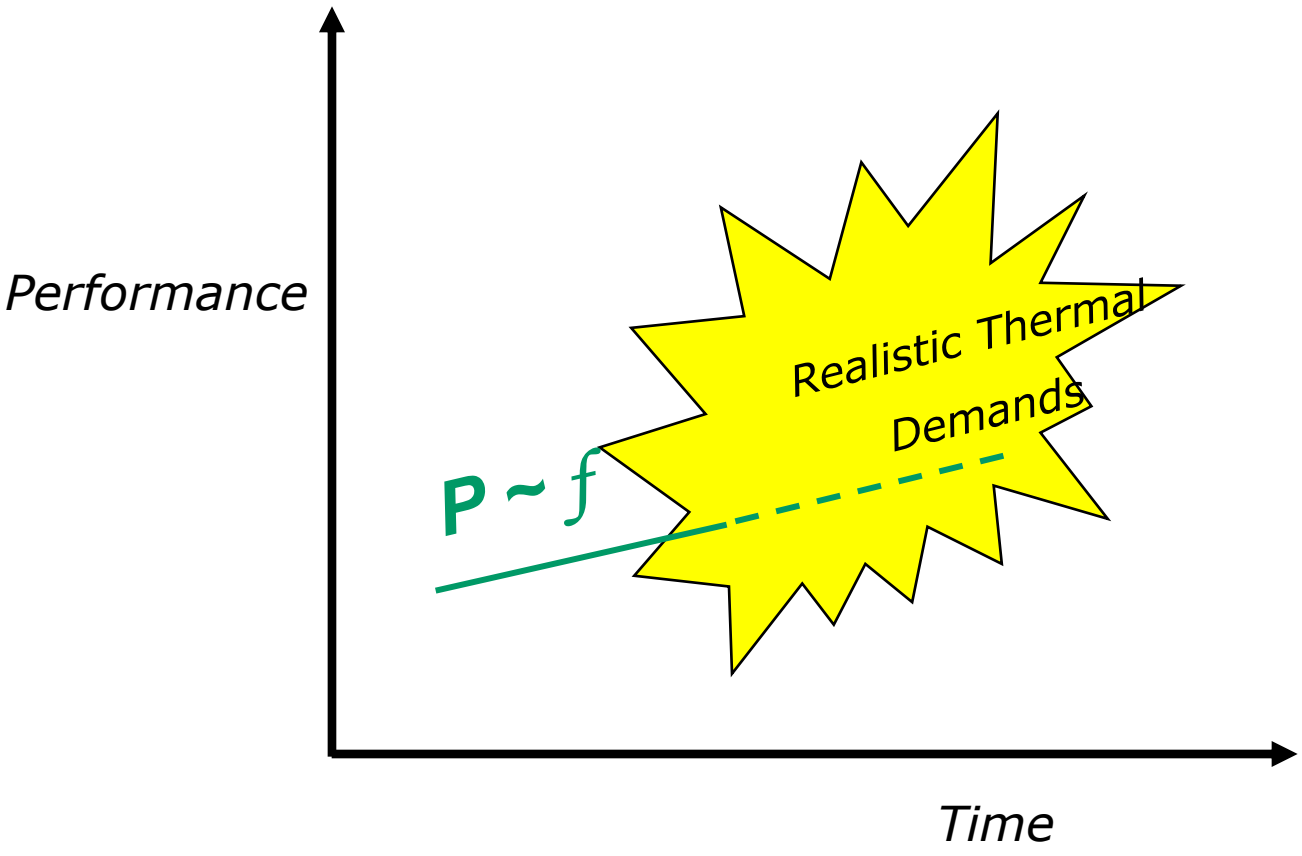

New Metric: Performance per Watt

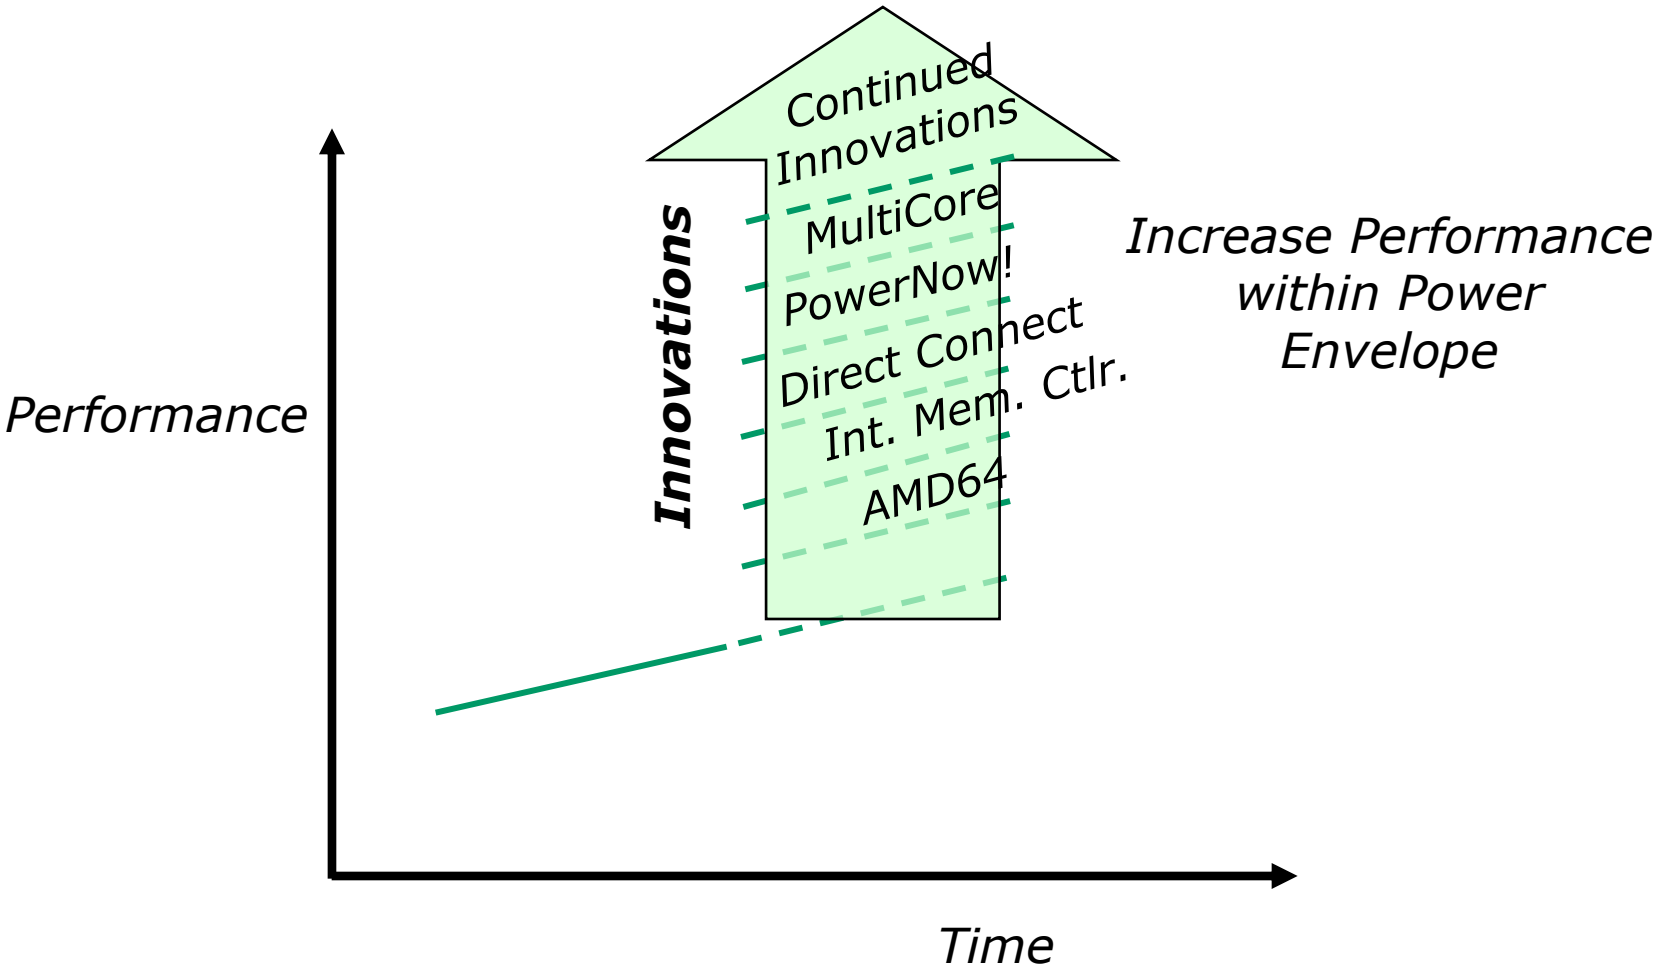

Power – The final frontier…

AMD
Smarter Choice

The combined total of data centers in California for 2004 were estimated to require ~300MW of energy.

That's equivalent to ~5000 barrels of oil a day!

SOURCE: California Energy Commission
http://www.energy.ca.gov/reports/2004-04-07_500-04-004.PDF

Feeding the Beast (Watts)

AMD Smarter Choice

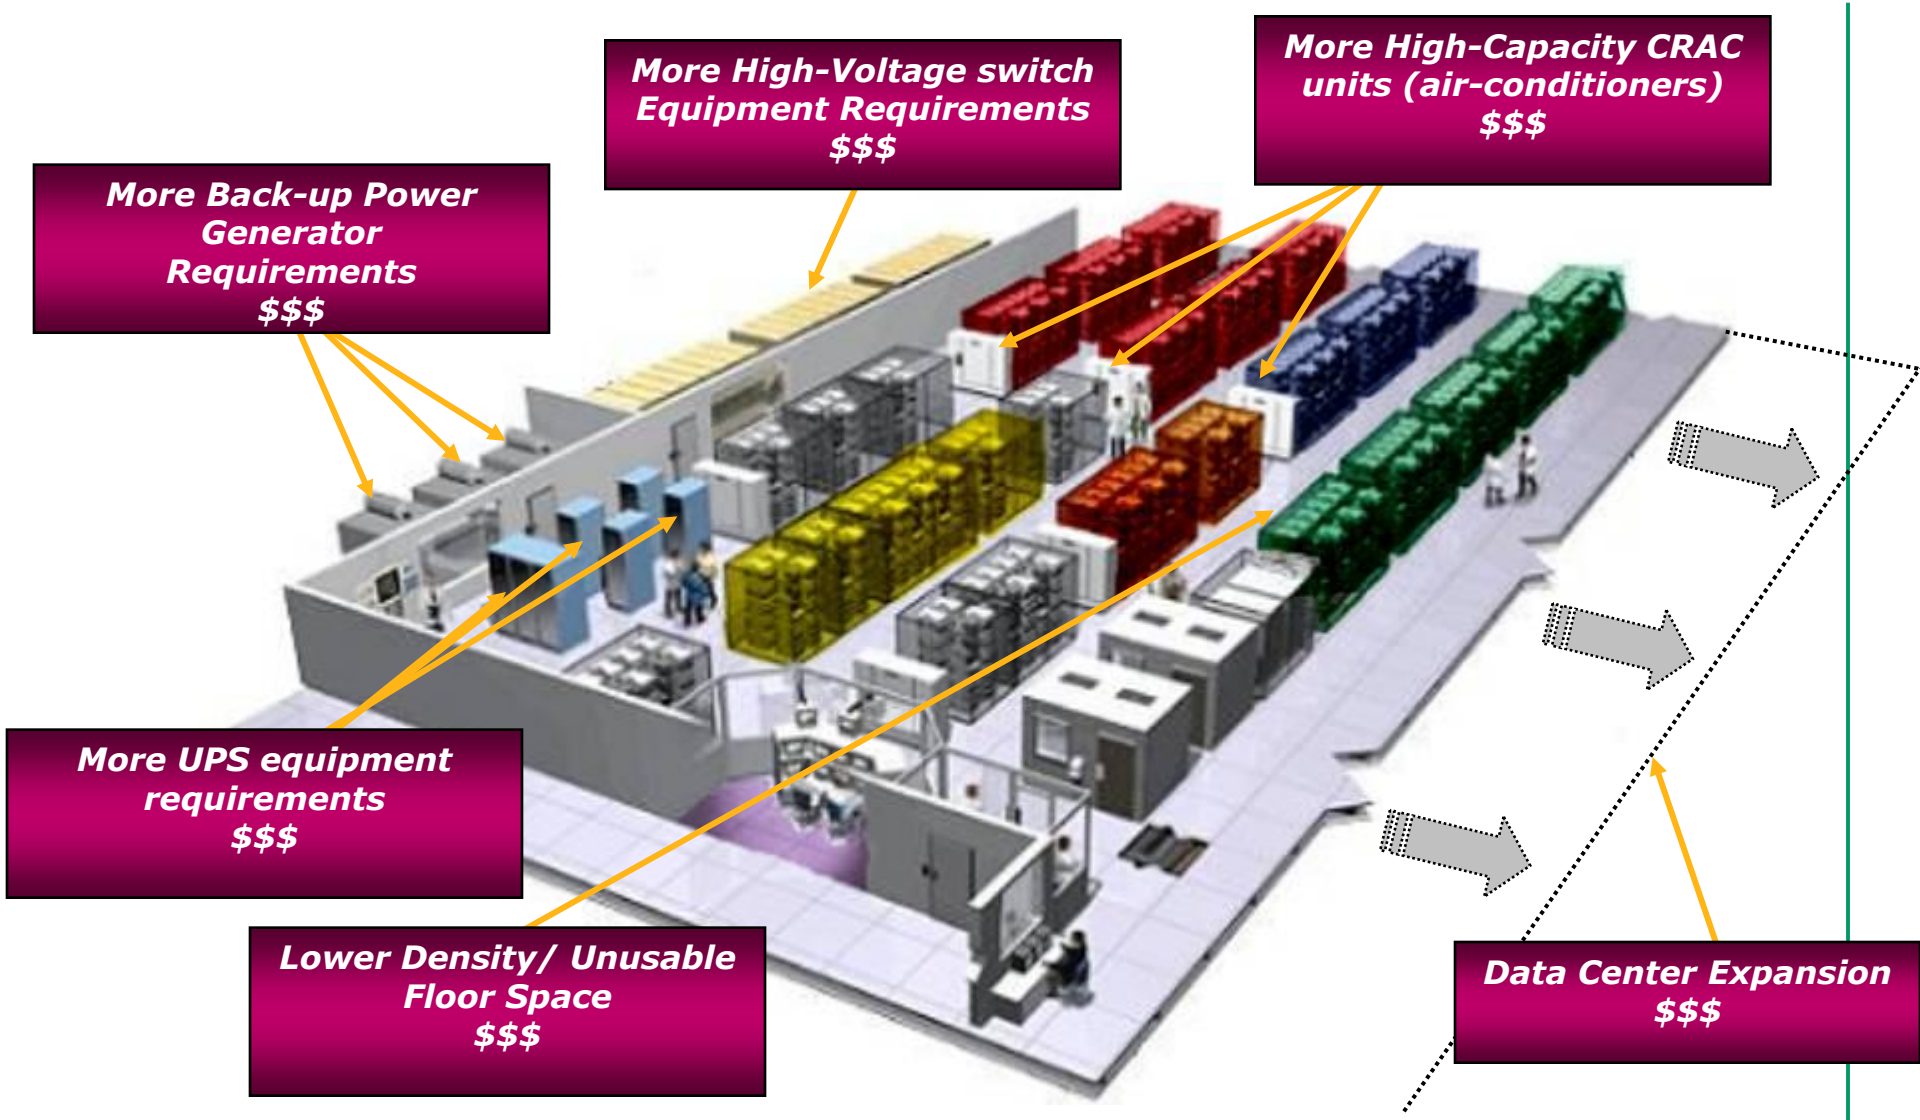

IT Knows it has a Problem

AMD Smarter Choice

Power Consumption/Cooling Issues Tracked By Company: 71%

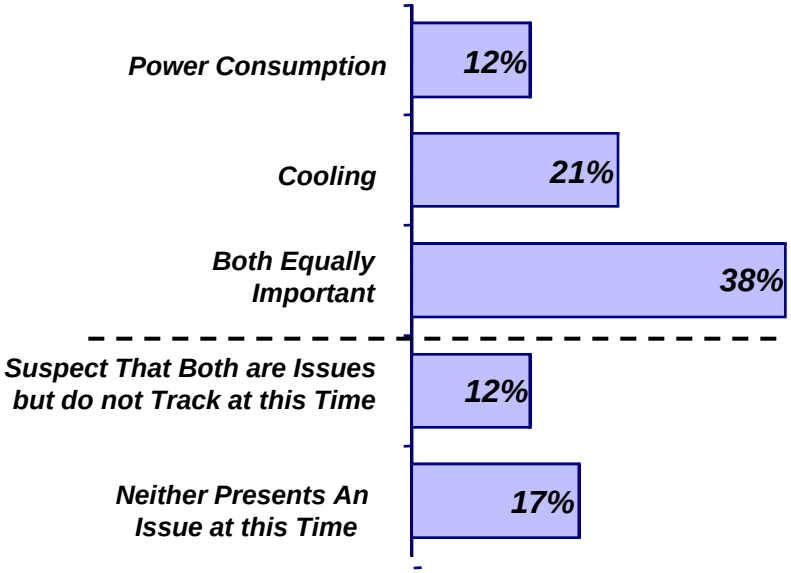

How Are Companies Addressing These Issues

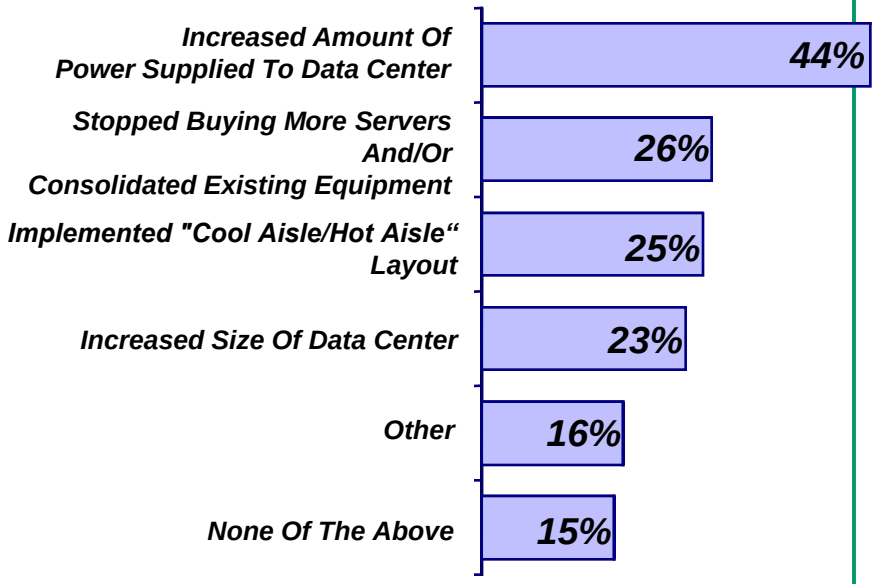

Base: 1,177 IT Decision Makers November 2005

Strategy Group/Ziff-Davis

*Average of **18% of total rack space wasted** due to power and cooling issues*

Next-Generation Power Comparison

In 2006 Next-Generation AMD Opteron™ Defined **A New Standard In Performance-Per-Watt** With Energy-Efficient **DDR2 Memory** and **Improved AMD PowerNow!™** Capabilities

In mid-2007 We Plan to Offer **Quad-Core AMD Opteron** in the Same DDR2-based Platforms at the **Same Power Efficiency**

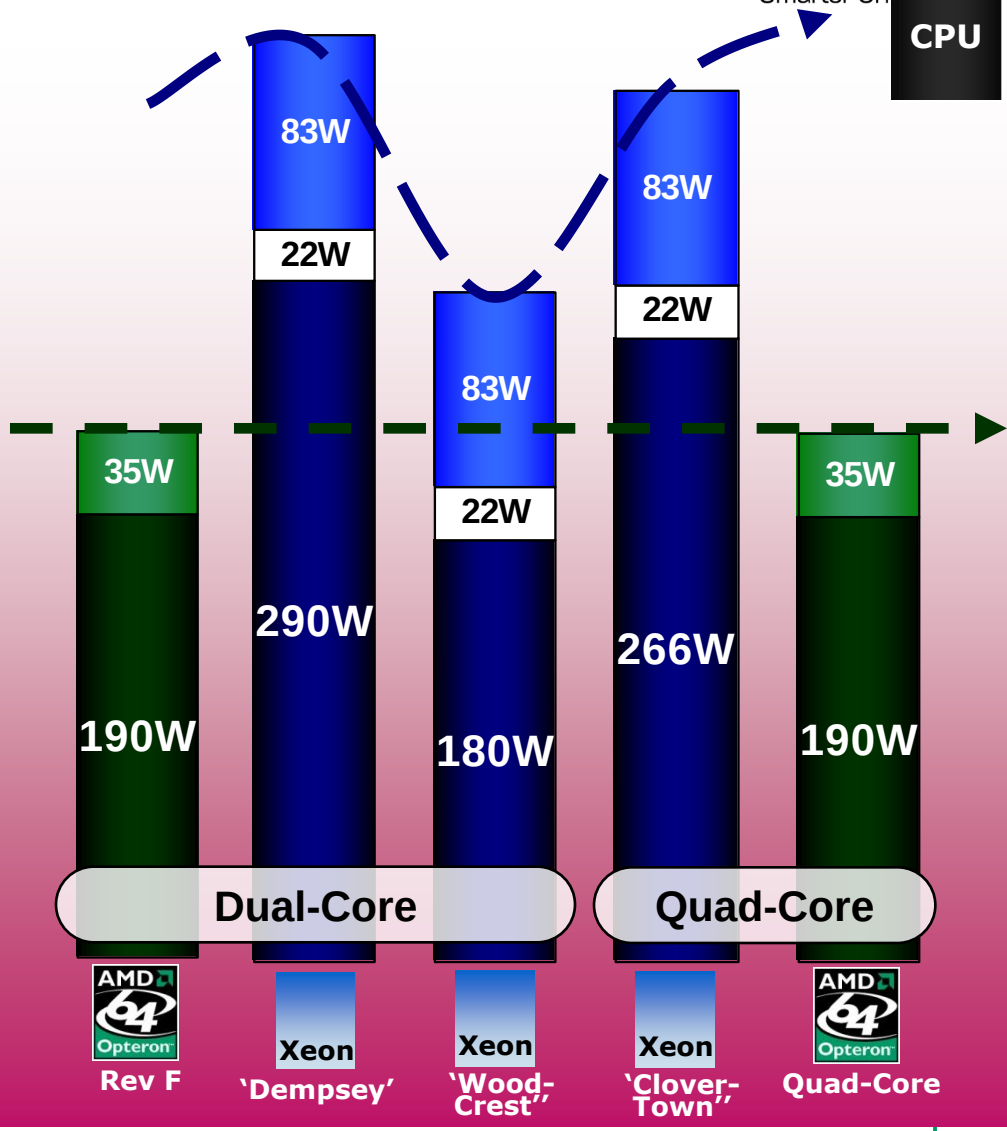

http://www.reghardware.co.uk/2006/05/25/intel_clovertown_power_specs/) and is subject to change. The examples contained herein are intended for informational purposes only. Other factors will affect real-world power consumption.

At the Wall Power Comparison

AMD Power Efficiency Innovation

AMD
Smarter Choice

Independent Dynamic Core Technology

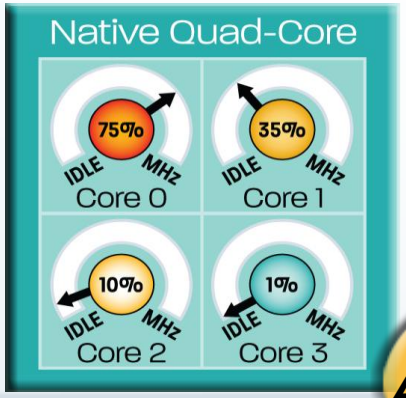

AMD CoolCore™ Technology

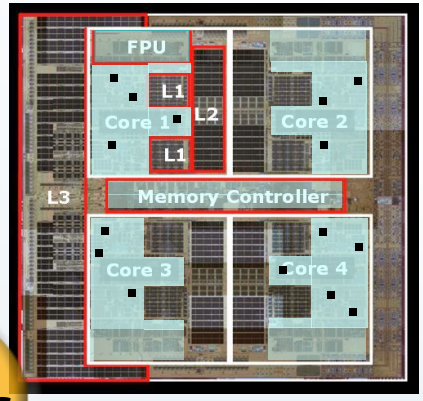

Same Power And Thermal Envelopes As Dual-Core!

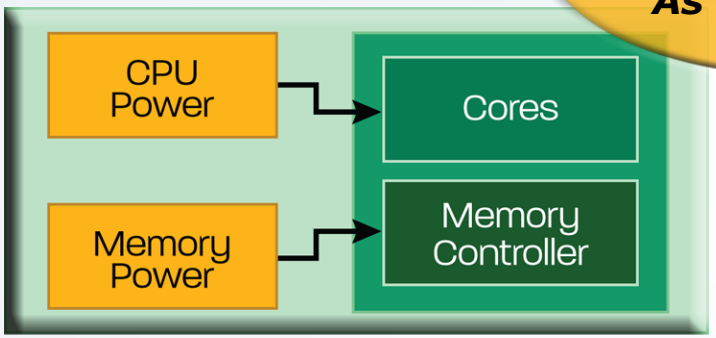

Dual Dynamic Power Management™

Low-Power DDR2 Memory

Computational Requirements

AMD Smarter Choice

Sun microsystems

What makes a great HPC CPU?

- Great CPU performance, FP and Integer
  - Lots of Cores at good clock rate
  - Maximum throughput per power
- Maximum Cache and Memory bandwidth
  - Each Flop is 2 Reads and 1 Write
  - Ideally 1 Byte/sec per FLOP
- Memory bandwidth must scale with # of cores
  - Otherwise the efficiency of additional cores is low
  - Major challenge beyond 4 cores per socket

A balanced system means to have good Memory and I/O Bandwidth / FLOP

Source: Andy Bechtolsheim , IEEE Grid & Cluster Conference, 2007

What Makes a Great CPU even better?

AMD Smarter Choice

- Frequency Lift

- Instruction Set Enhancement

- Increasing Cores

- Core Capabilities

- Memory Access (Capacity, Bandwidth & Latency)

- CPU-CPU connectivity

- IO Connectivity (Data access)

- Software Scaling

Legacy Northbridge Server Architecture

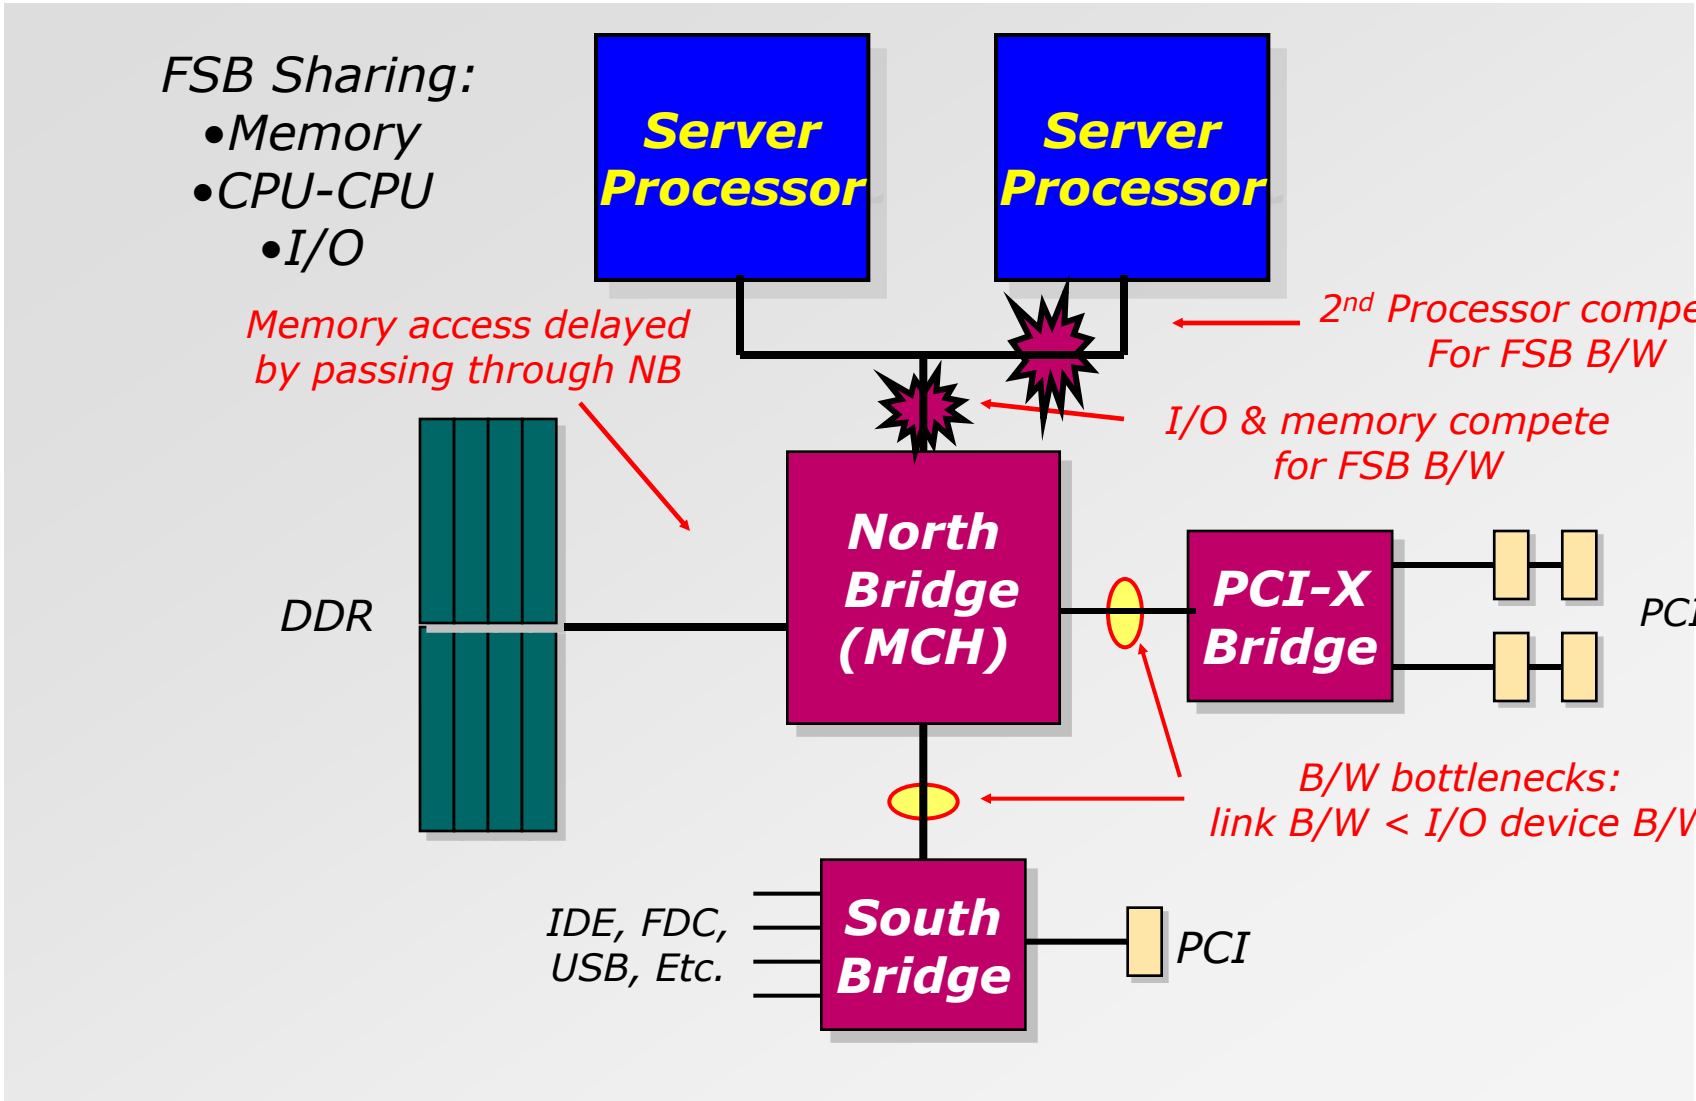

AMD Opteron™ "Direct Connect" Server Architecture

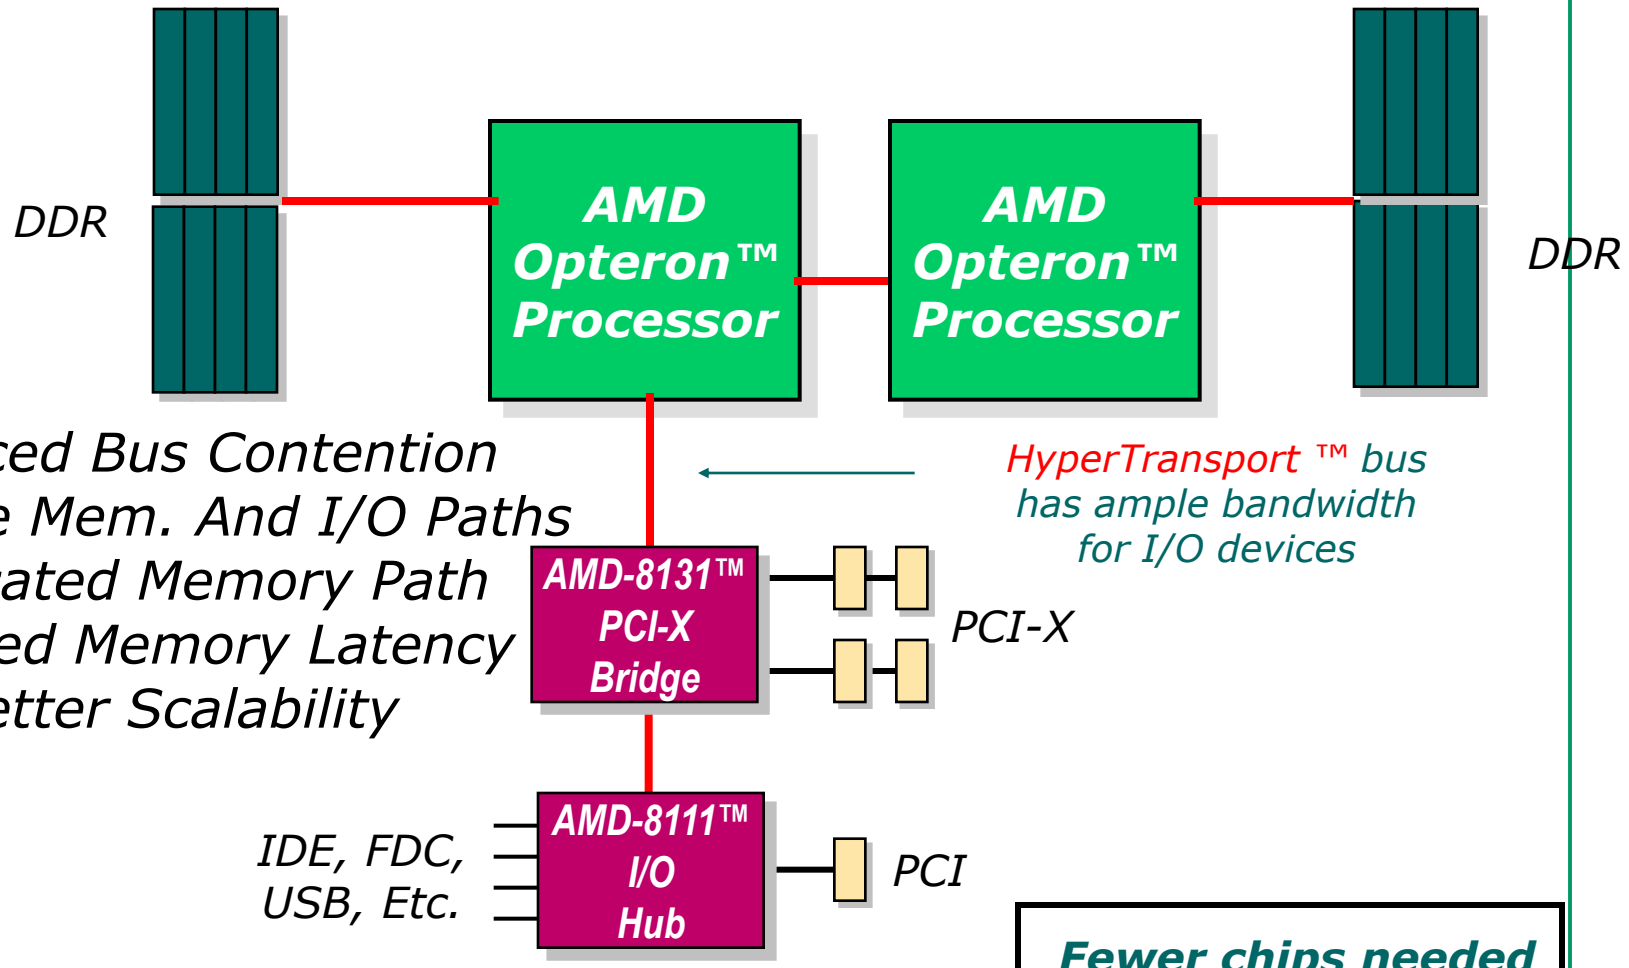

- Reduced Bus Contention
- Separate Mem. And I/O Paths
- Dedicated Memory Path
- Reduced Memory Latency
- Better Scalability

Fewer chips needed for basic server

AMD Opteron™ Platform

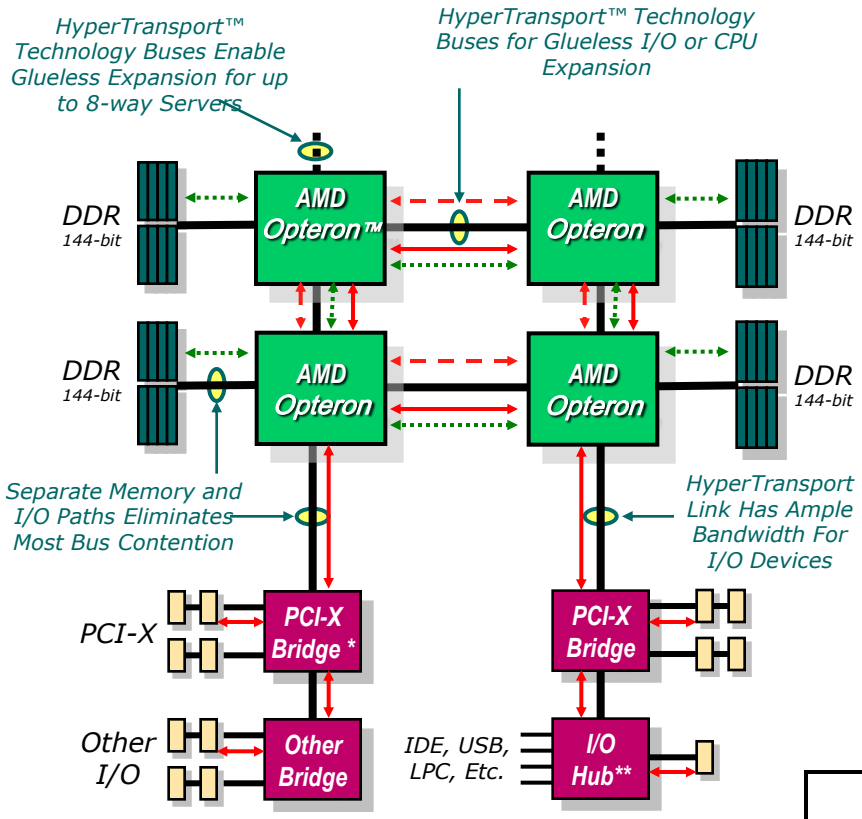

Fewer Chips Needed for 4-way Server (Reduces Cost)

Historic MP Server

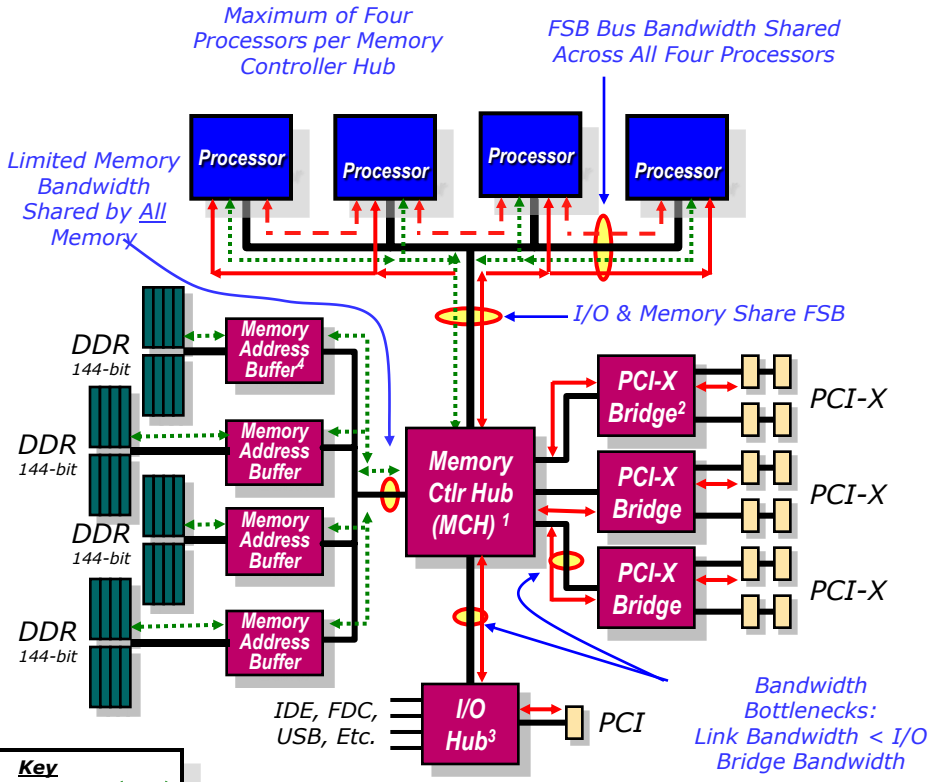

Key
Memory Traffic
I/O Traffic
IPC Traffic

9-Chip Chipset Needed for 4-way Server

- **Scalable memory and I/O bandwidth**
  - Up to 8 processors without glue logic
  - Each processor adds more memory
  - Each processor adds additional HyperTransport™ buses for more PCI-X and other I/O bridges
  - Fewer chips required

*AMD-8131™ HyperTransport PCI-X Tunnel **AMD-8111™ HyperTransport I/O Hub

- **System scalability limited by Northbridge**
  - Maximum of 4 processors
    - Processors compete for FSB bandwidth
  - Memory size and bandwidth are limited
  - Maximum of 3 PCI-X bridges
  - Many more chips required

[1] ServerWorks CMIC HE Memory Controller Hub (MCH) [2] ServerWorks CIOB-X 64-bit PCI/PCI-X Controller Hub
[3] ServerWorks CSB5 I/O Controller Hub [4] ServerWorks REMC Memory Address Buffer

AMD Performance Innovation

AMD
Smarter Choice

AMD Wide Floating-Point Accelerator

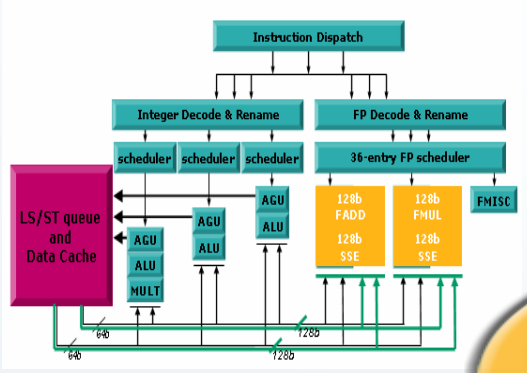

AMD Memory Optimizer Technology

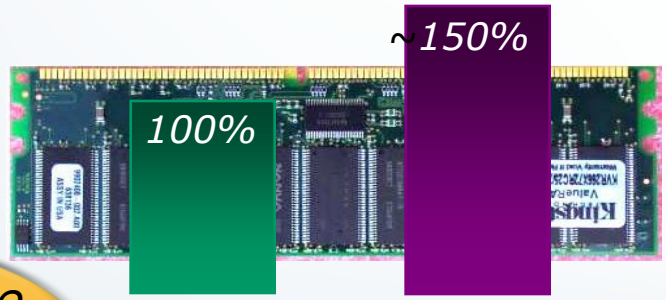

Comprehensive Performance Enhancements!

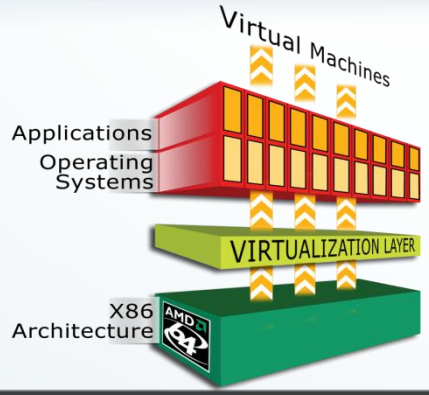

Rapid Virtualization Indexing™

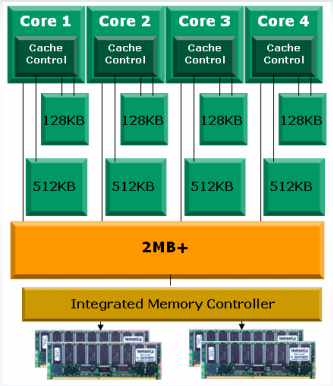

AMD Balanced Smart Cache

Dual-Core to Quad-Core Uplift

AMD Smarter Choice

Dual-Core AMD Opteron™ 2200 Series vs. Quad-Core AMD Opteron Model 2350 2 Socket Performance Scaling

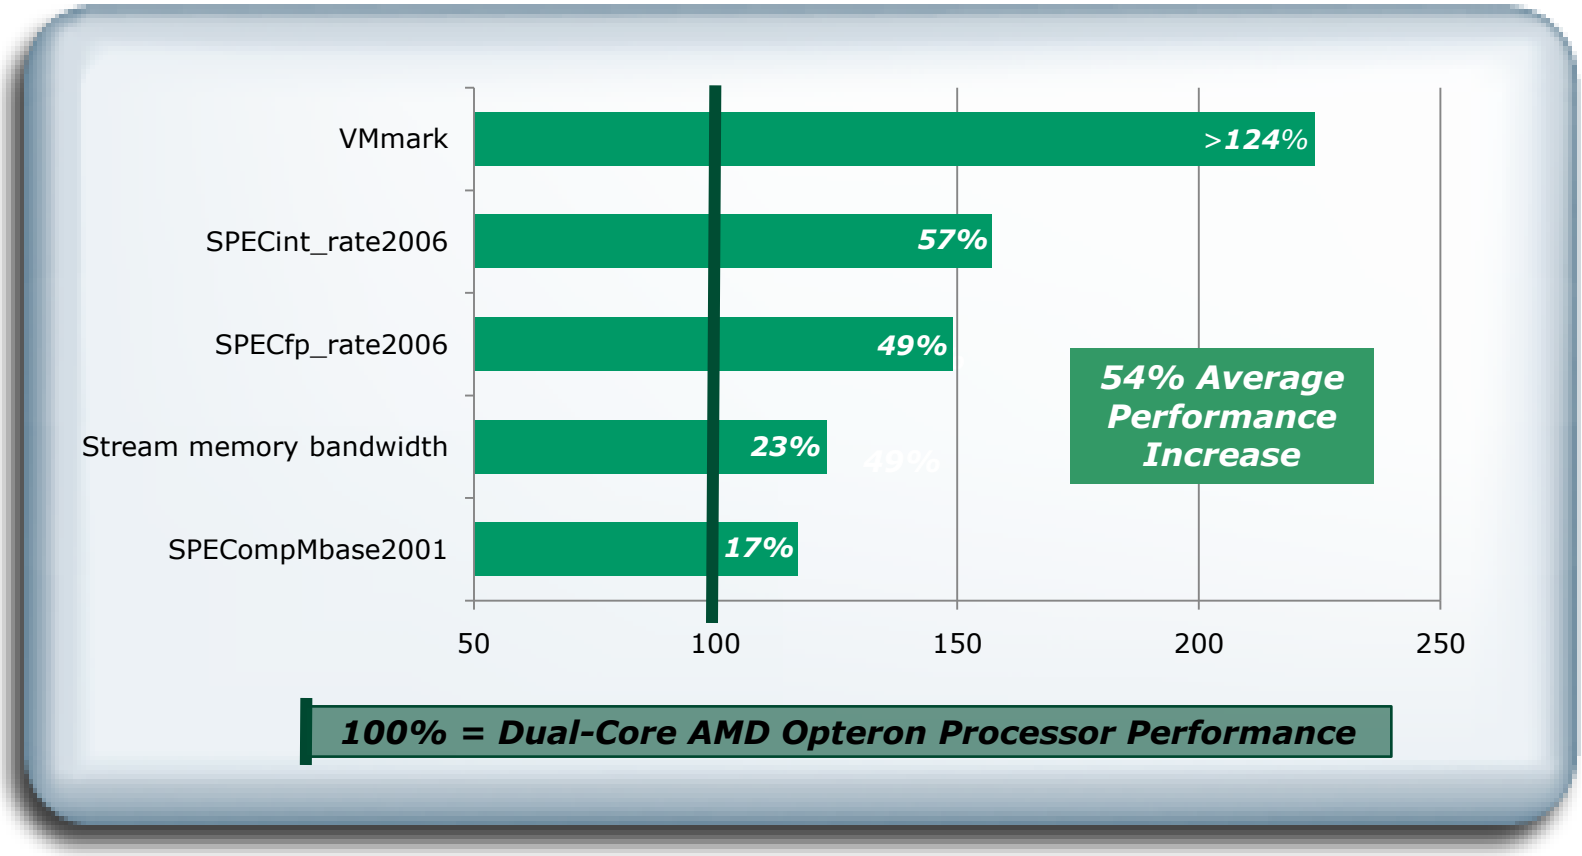

SPEC and the benchmark name SPECint, SPECfp and SPECOMPM are registered trademarks of the Standard Performance Evaluation Corporation. Benchmark results stated above for Dual-Core AMD Opteron™ processor Model 2222 reflect results published on www.spec.org as of Sep 9, 2007. The comparison presented above is based on results for Quad-Core AMD Opteron processor Model 2350 under submission to SPEC as of Sep 9, 2007. For the latest results visit http://www.spec.org/cpu2006/results/ and http://www.spec.org/omp/results/. Stream and VMmark results based on internal measurements at AMD performance labs.

Quad-Core vs. Dual-Die

AMD’s design will be a TRUE quad-core processor without compromising performance, power or heat

Intel may rush a “dual die” architecture to market in order to claim “first to market”, only to change the design to true quad-core later - more churn and increased customer TCO

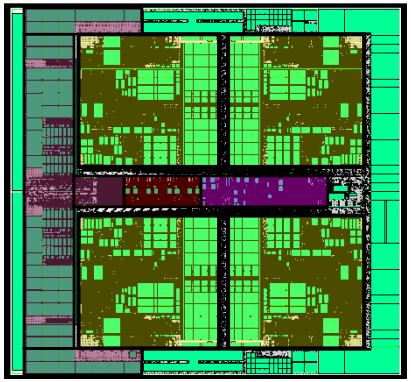

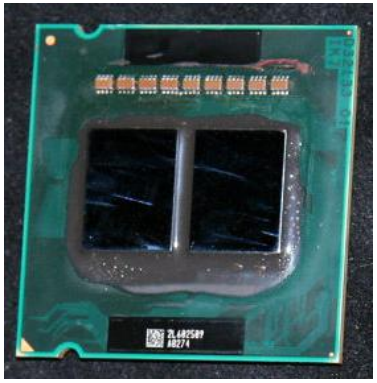

Native Quad-Core Design
- Optimum performance
- Same power & thermal envelopes as dual-core

Dual Die (Dual Cavity)
- Can hinder performance (FSB design)
- Publicly known thermal design power ranges higher than dual-core products

Native Quad-Core Benefit:
Faster Data Sharing

AMD Smarter Choice

Situation: Core 1 needs data in Core 3 cache … How Does it Get There?

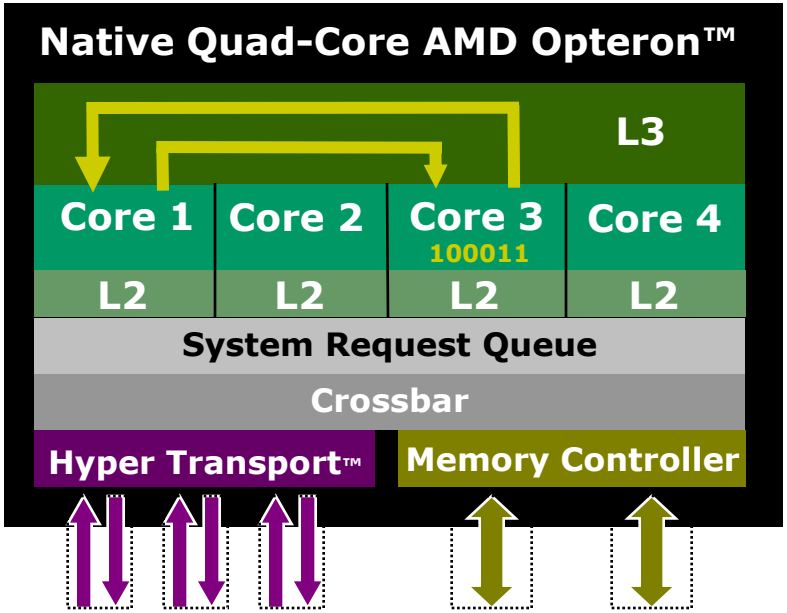

1. Core 1 probes Core 3 cache, data is copied directly back to Core 1

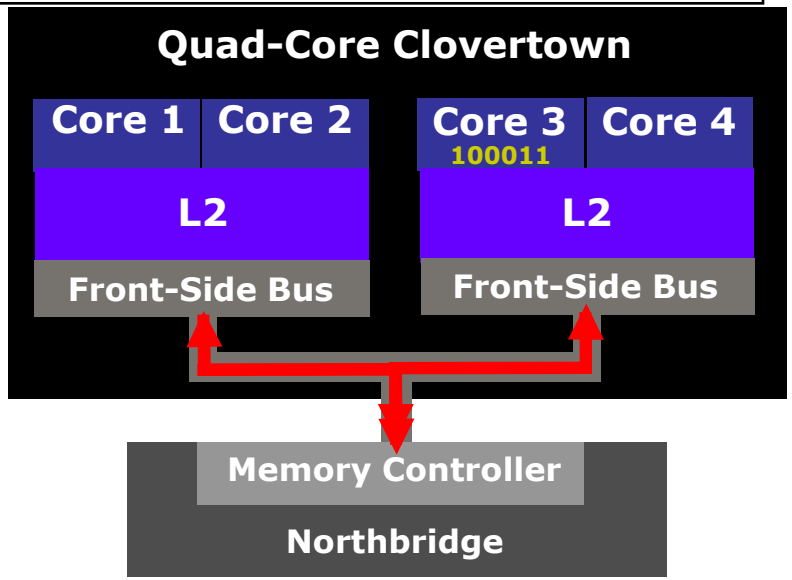

1. Core 1 sends a request to the memory controller, which probes Core 3 cache
2. Core 3 sends data back to the memory controller, which forwards it to Core 1

This happens at processor frequency

Result: Improved Quad-Core Performance

This happens at front-side bus frequency

Result: Reduced Quad-Core Performance

Performance-Per-Watt Leadership

AMD Smarter Choice

Quad-Core AMD Opteron™ Processor Model 2350 (75 Watt) vs. Intel Xeon 5345 (80 Watt, without Additional Watts of Memory Controller and FBDIMM)

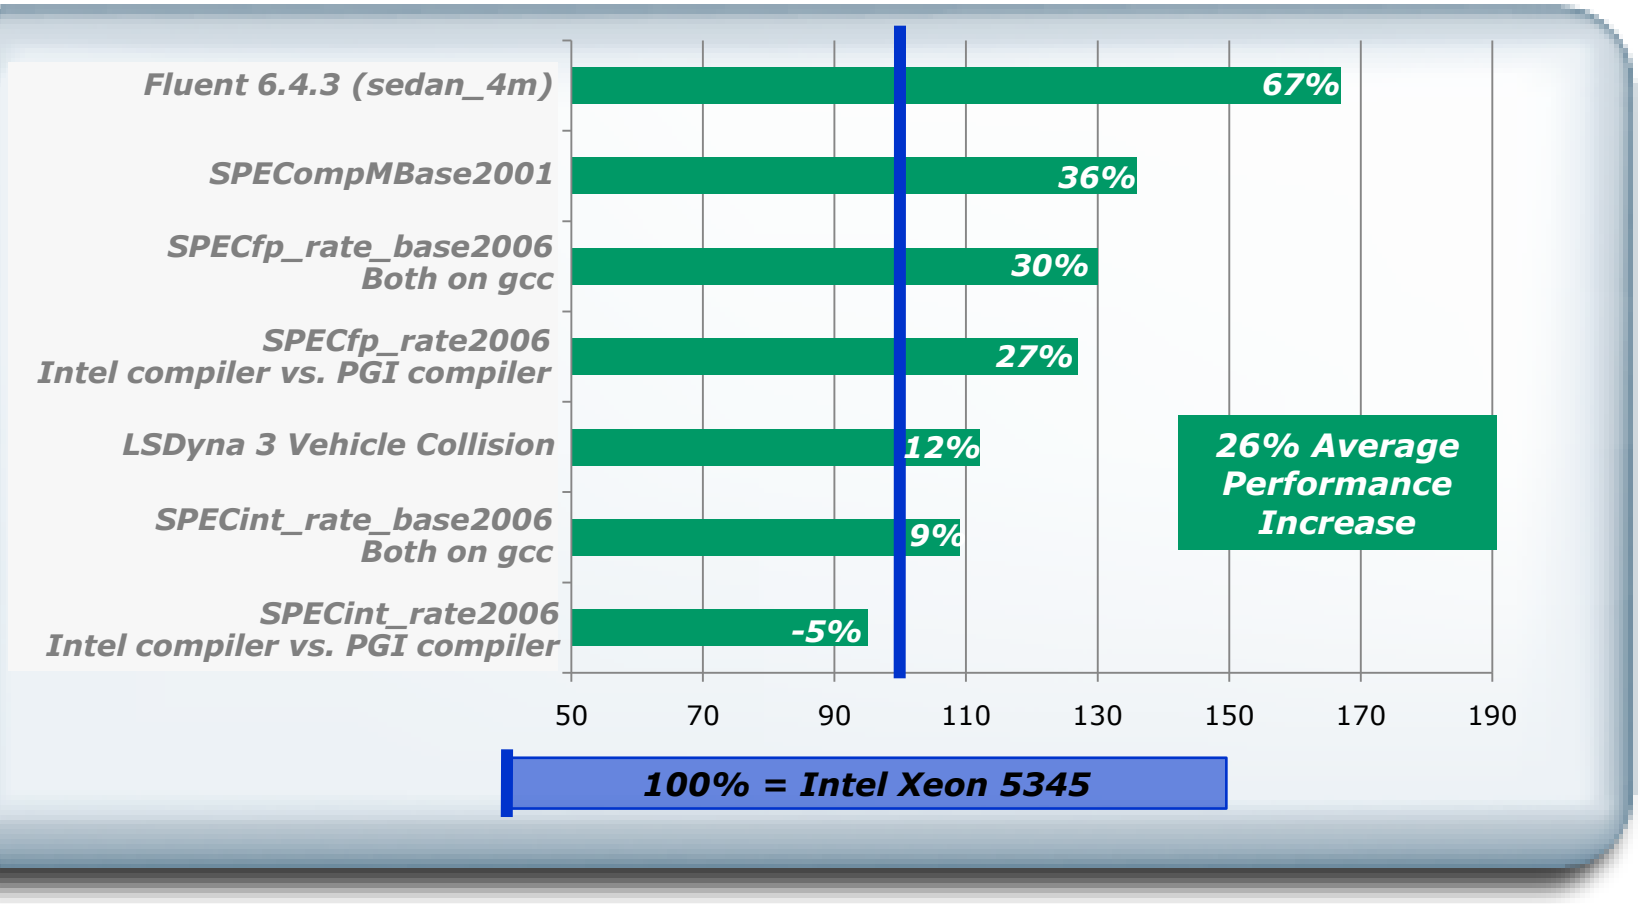

SPEC and the benchmark name SPECint, SPECfp and SPECOMPM are registered trademarks of the Standard Performance Evaluation Corporation. Competitive benchmark results stated above reflect results published on www.spec.org as of Sep 9, 2007. The comparison presented above is based on results for Quad-Core AMD Opteron processor Model 2350 and Xeon 5345 (specint_rate2006 gcc and SPECompM2001 base) under submission to SPEC as of Sep 9, 2007. For the latest results visit http://www.spec.org/cpu2006/results/. Fluent and LSDyna result based on internal measurements at AMD performance labs.

AMD Server Platform Roadmap

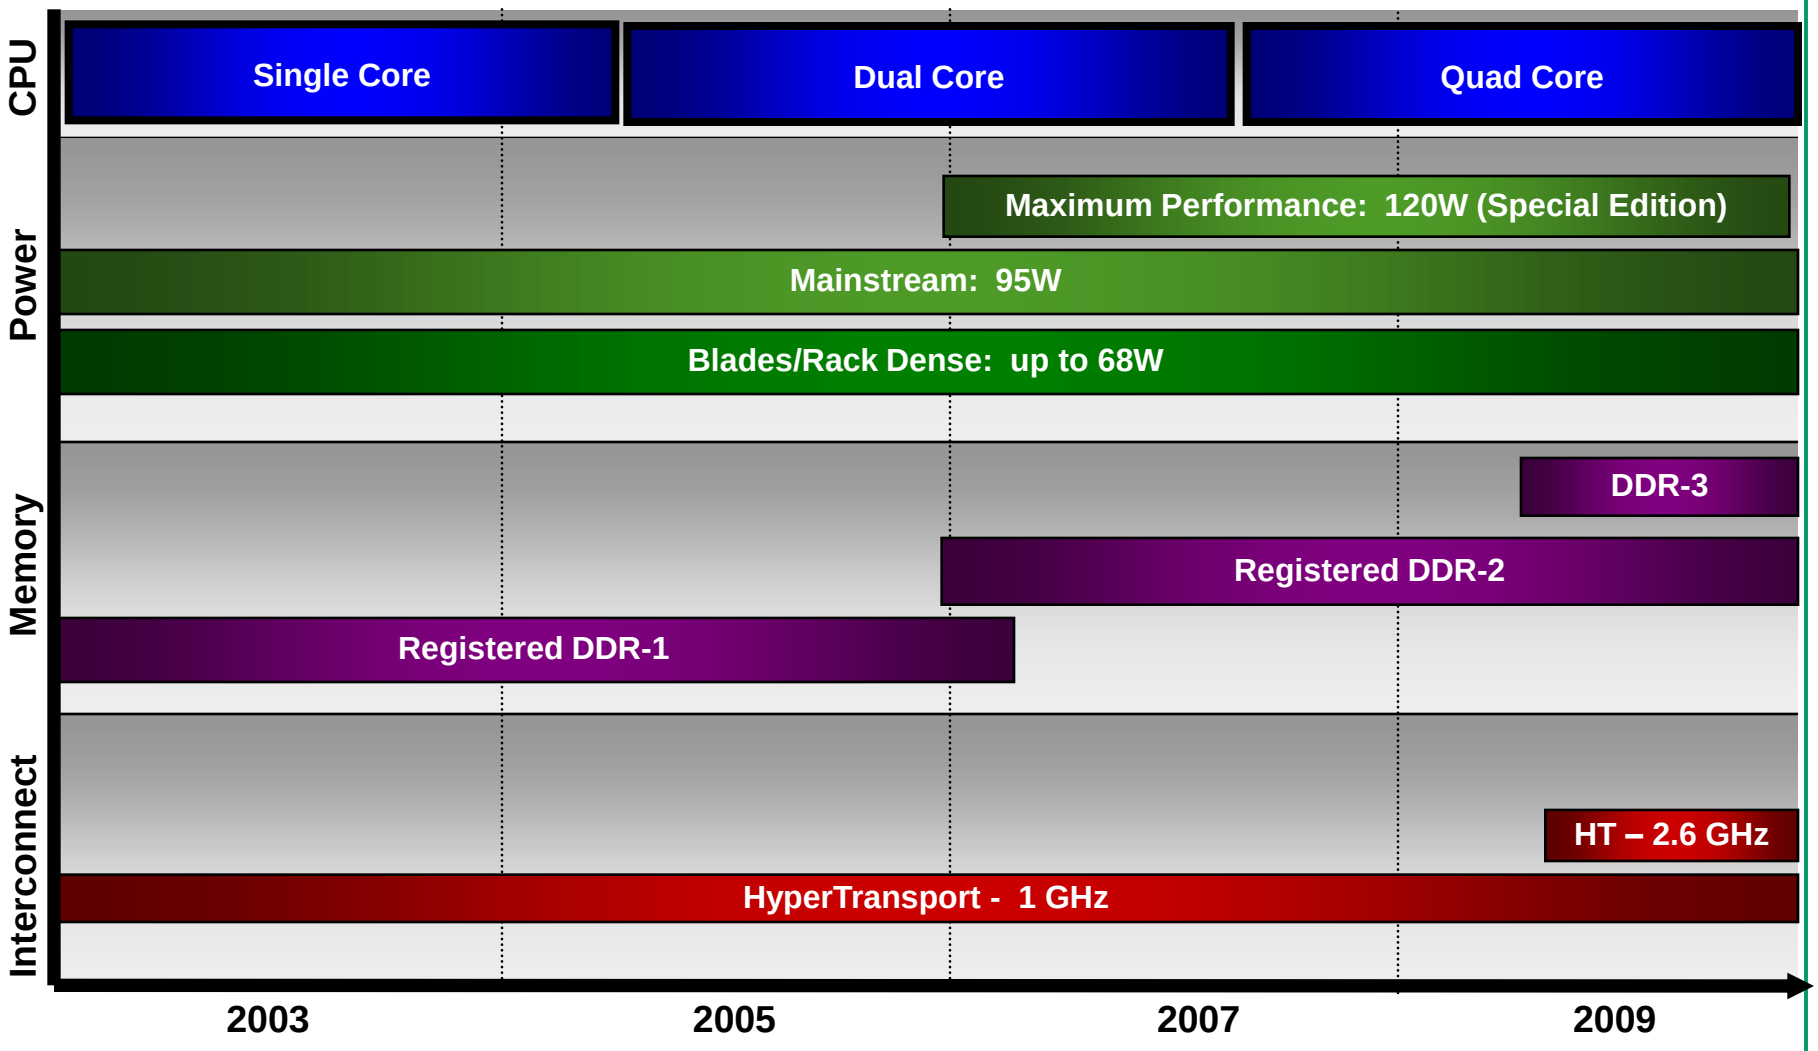

Investment Protection: Stable Platform Progression

Long-term success for partners and end-customers

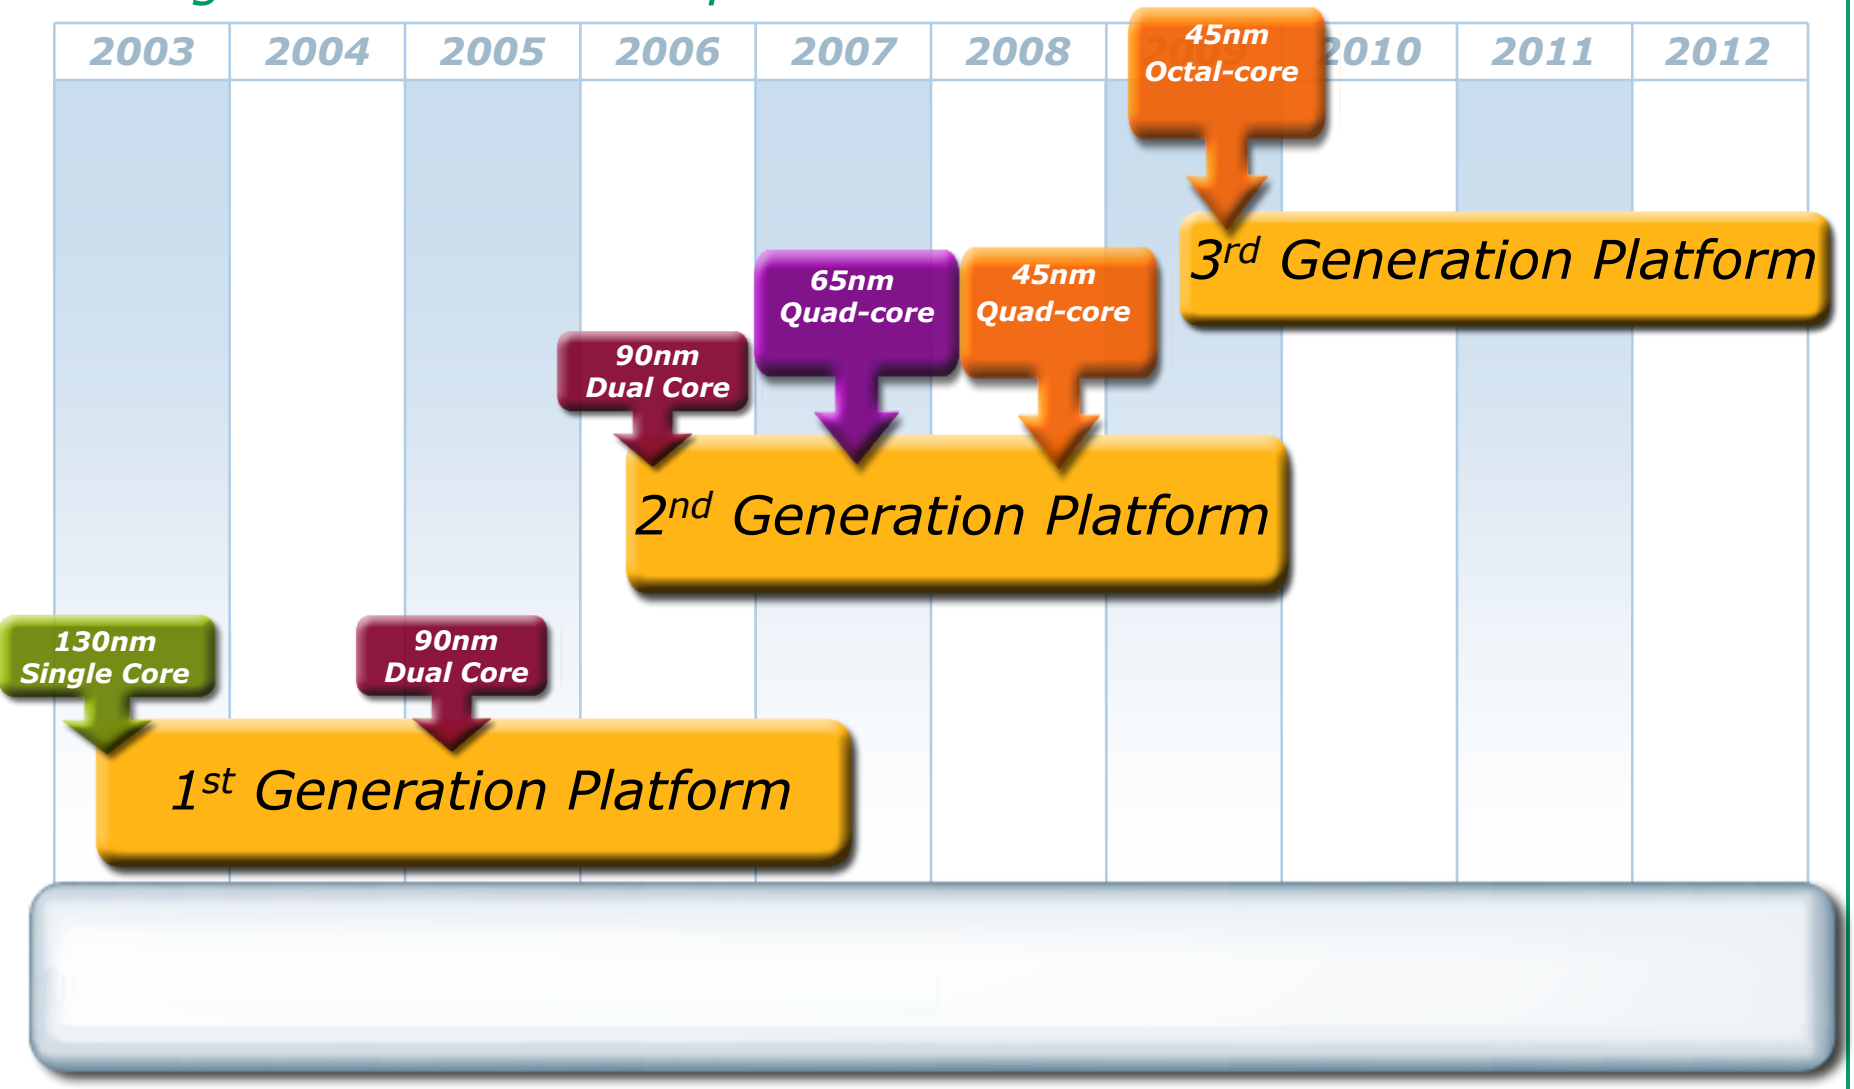

Platforms in Market Today

AMD Smarter Choice

2003 vs.

IBM eServer 325

1st Generation AMD Opteron™

2007

AMD Opteron™ 64

Sun microsystems
- X4600
- X2200
- X4500
- X4200/4100
- X2100
- U20 & U40 WS
- Blade X6220
- Blade X8420

DELL
- SC 1435
- PowerEdge 2970
- PowerEdge 6950

hp
- DL585
- DL385
- DL365
- DL145
- BL465c
- BL685c
- xw9400

Gateway
- E-9422R
- E-9522R
- E-9722R
- E-9222T

IBM
- X3455
- X3655
- X3755
- LS21
- LS41

CRAY
- XT3™
- XT4

egenera
- BladeFrame ES and EX

acer
- G5450

FUJITSU COMPUTERS SIEMENS
- X630 S2

AMD Validated Solutions

Customers are Responding!

AMD Opteron Server Market Share

AMD
Smarter Choice

	2005	2006
WW total	11.9%	16.0%
WW x86	12.8%	17.1%
WW x86 2-way	13.6%	18.2%
WW x86 4-way	28.2%	40.1%
US x86	20.8%	27.4%
US x86 2-way	21.2%	27.9%
US x86 4-way	36.0%	56.2%

Source: Gartner, IDC, end of year results

The Heterogeneous Processing Imperative

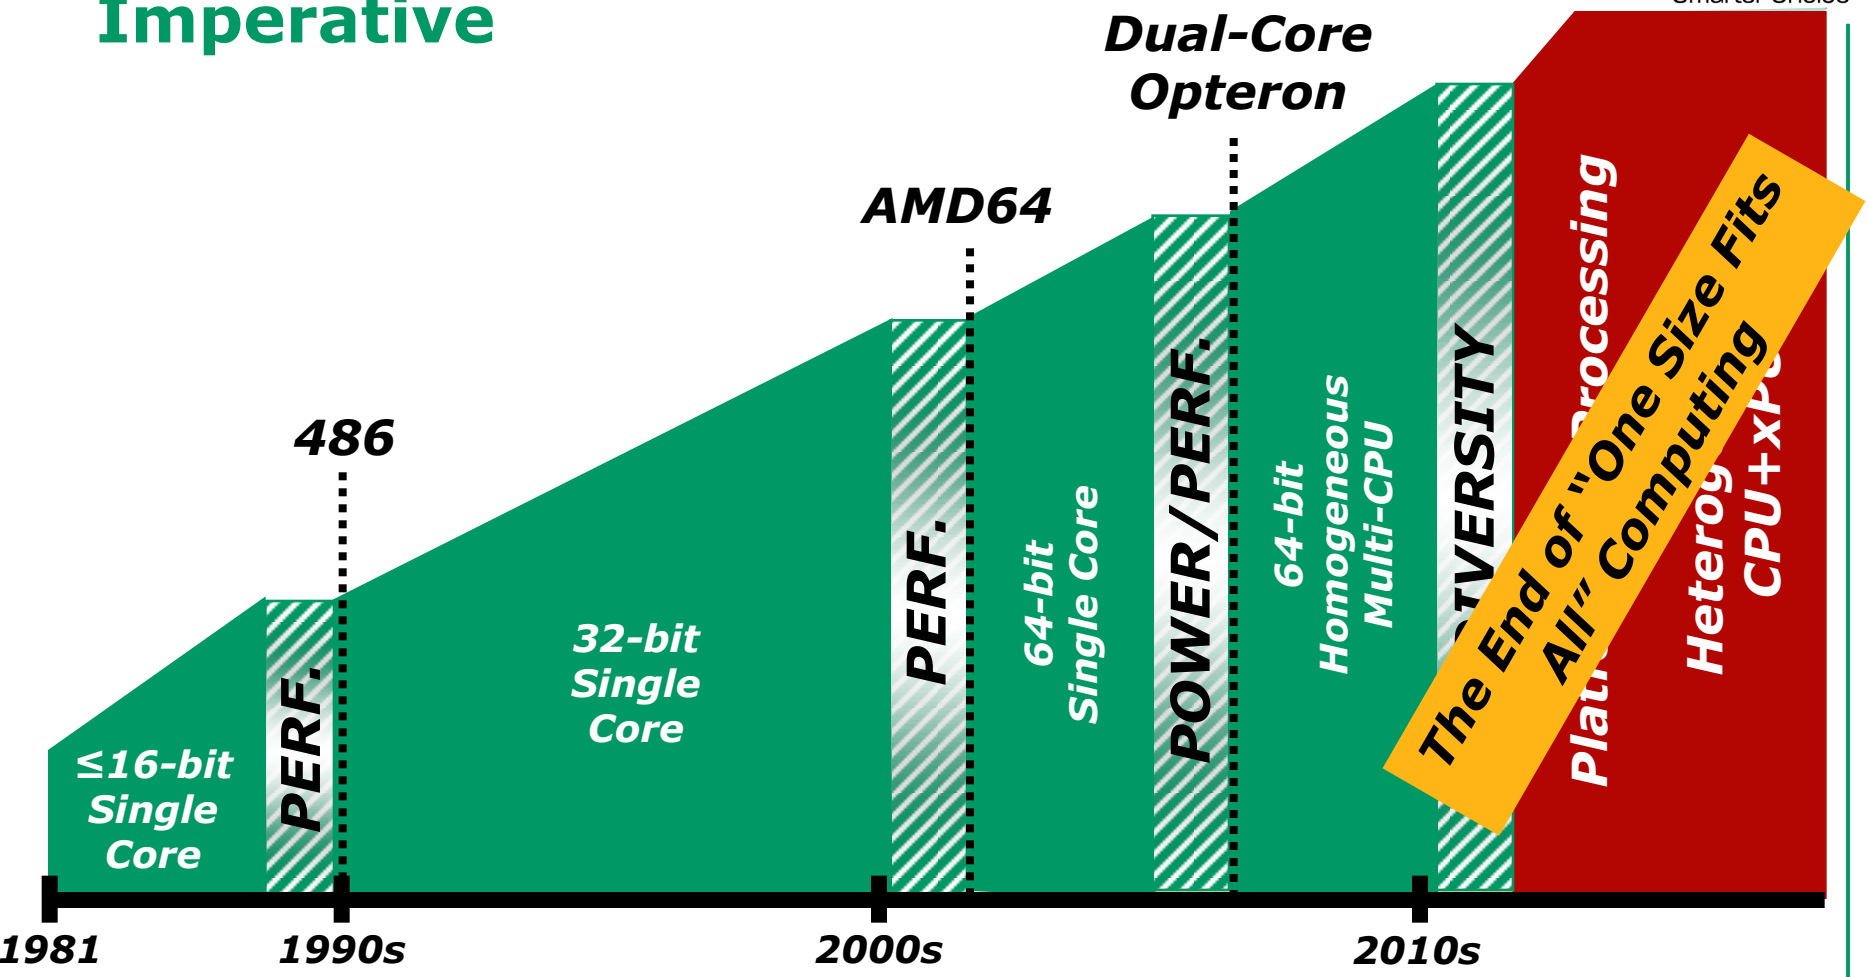

By the end of the decade, homogenous multi-core becomes increasingly inadequate

Continuum of Solutions

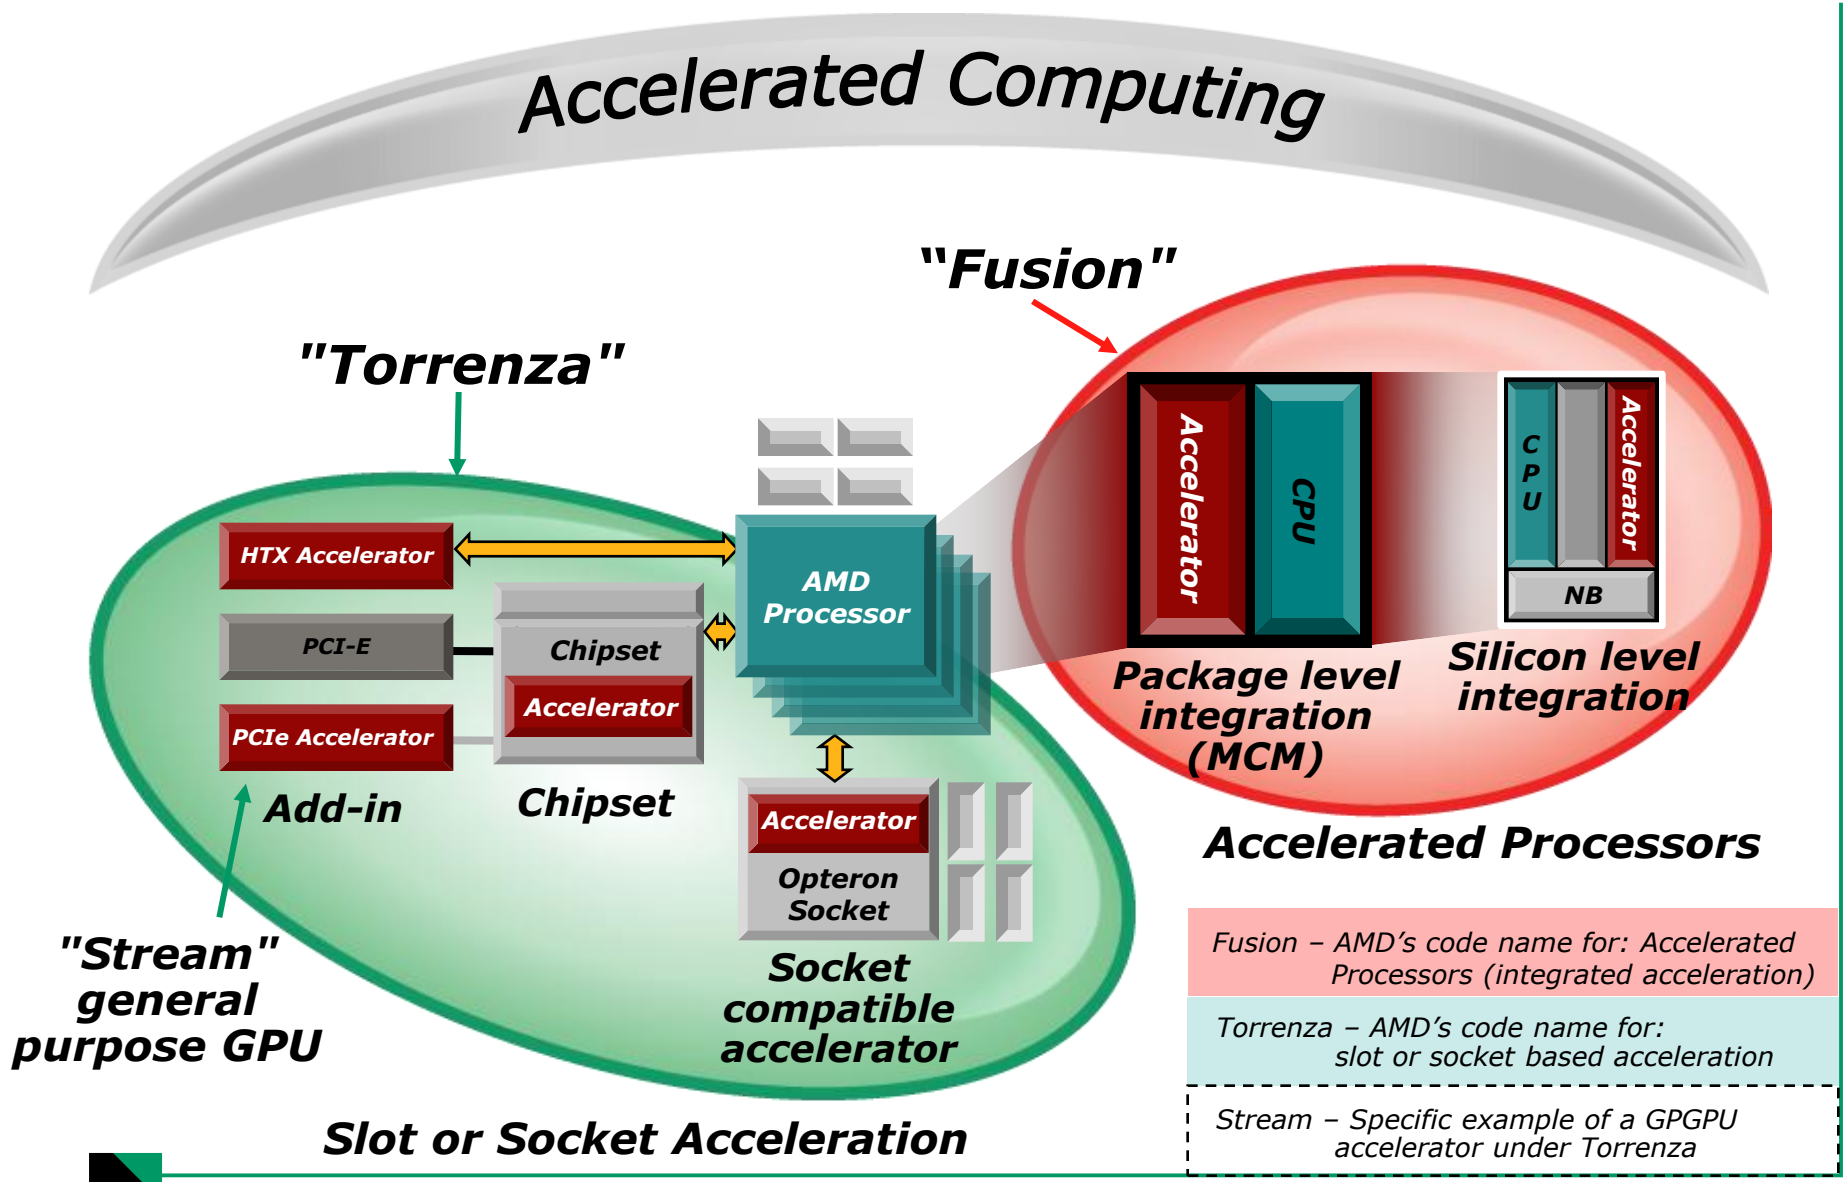

Fusion – AMD's code name for: Accelerated Processors (integrated acceleration)

Torrenza – AMD's code name for: slot or socket based acceleration

Stream – Specific example of a GPGPU accelerator under Torrenza

Torrenza: Accelerated Computing Today

AMD Smarter Choice

- *Direct Connect Accelerators in sockets or slots deliver superior performance without bridge chips*
- *100s of GFlops to solve complex math*

- *Familiar programming interfaces speed time to implementation*

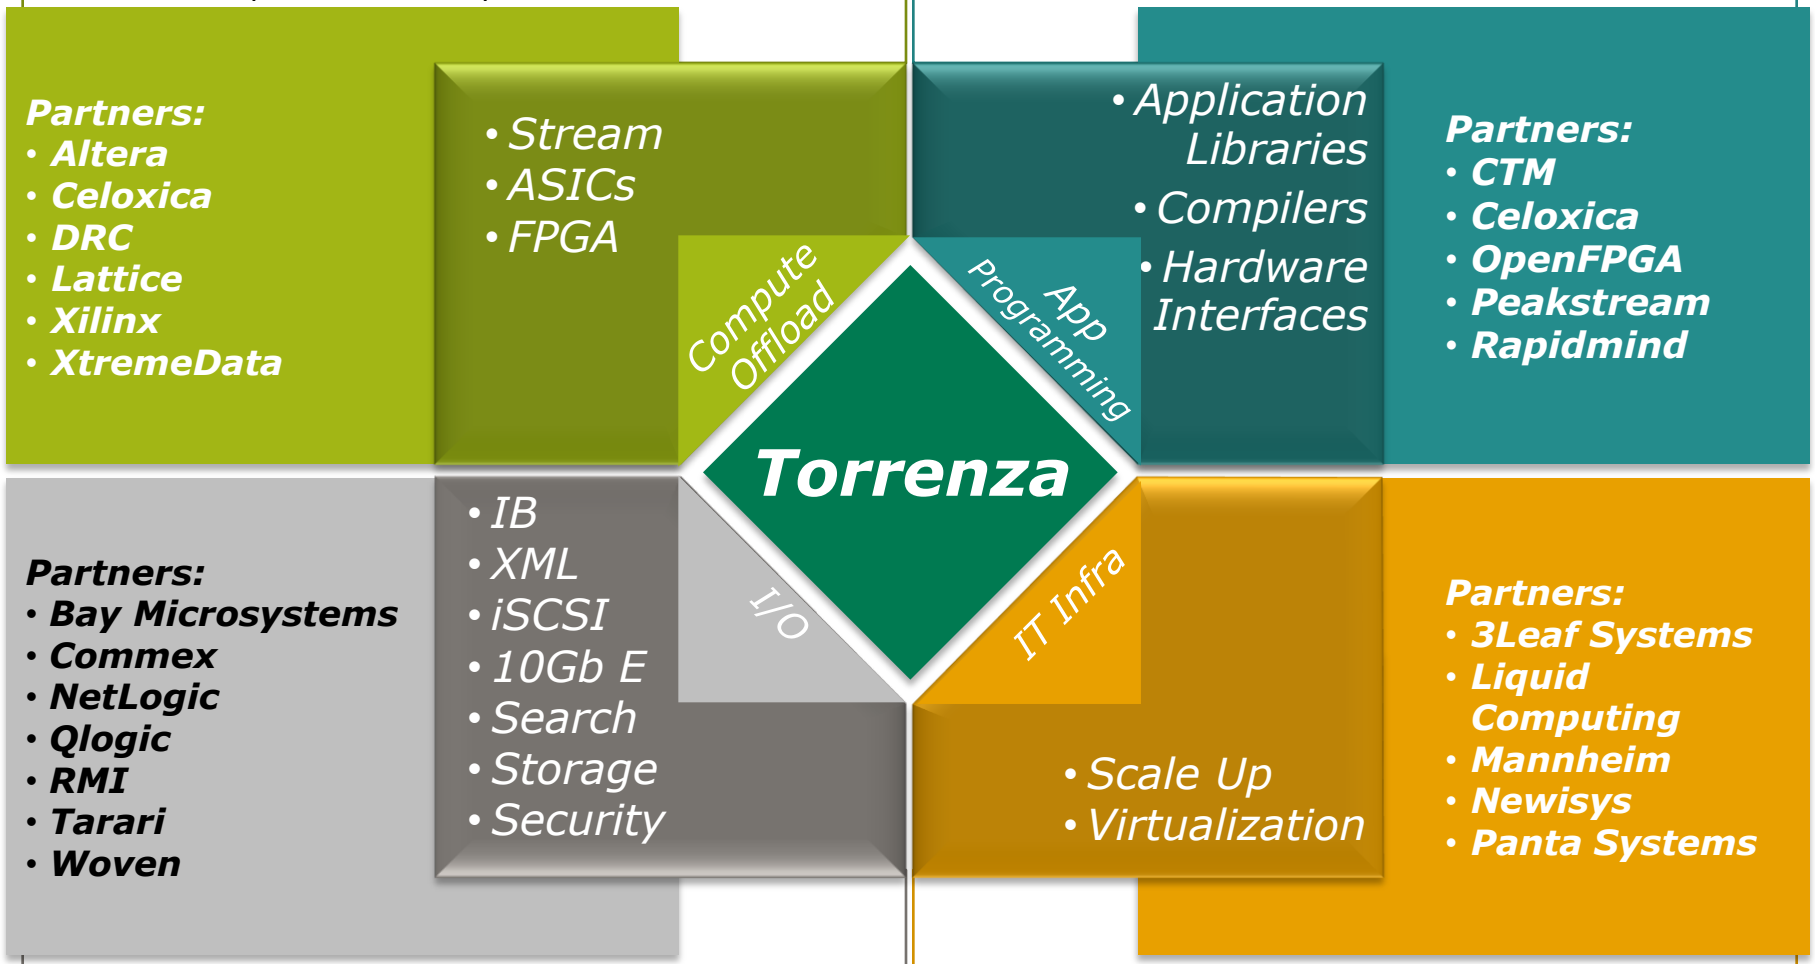

- *Specialized Direct Connect devices for high-throughput, low-latency processing*

- *cHT and HT provide peer level interfaces to build systems from commodity building blocks*

How Do We Get There?

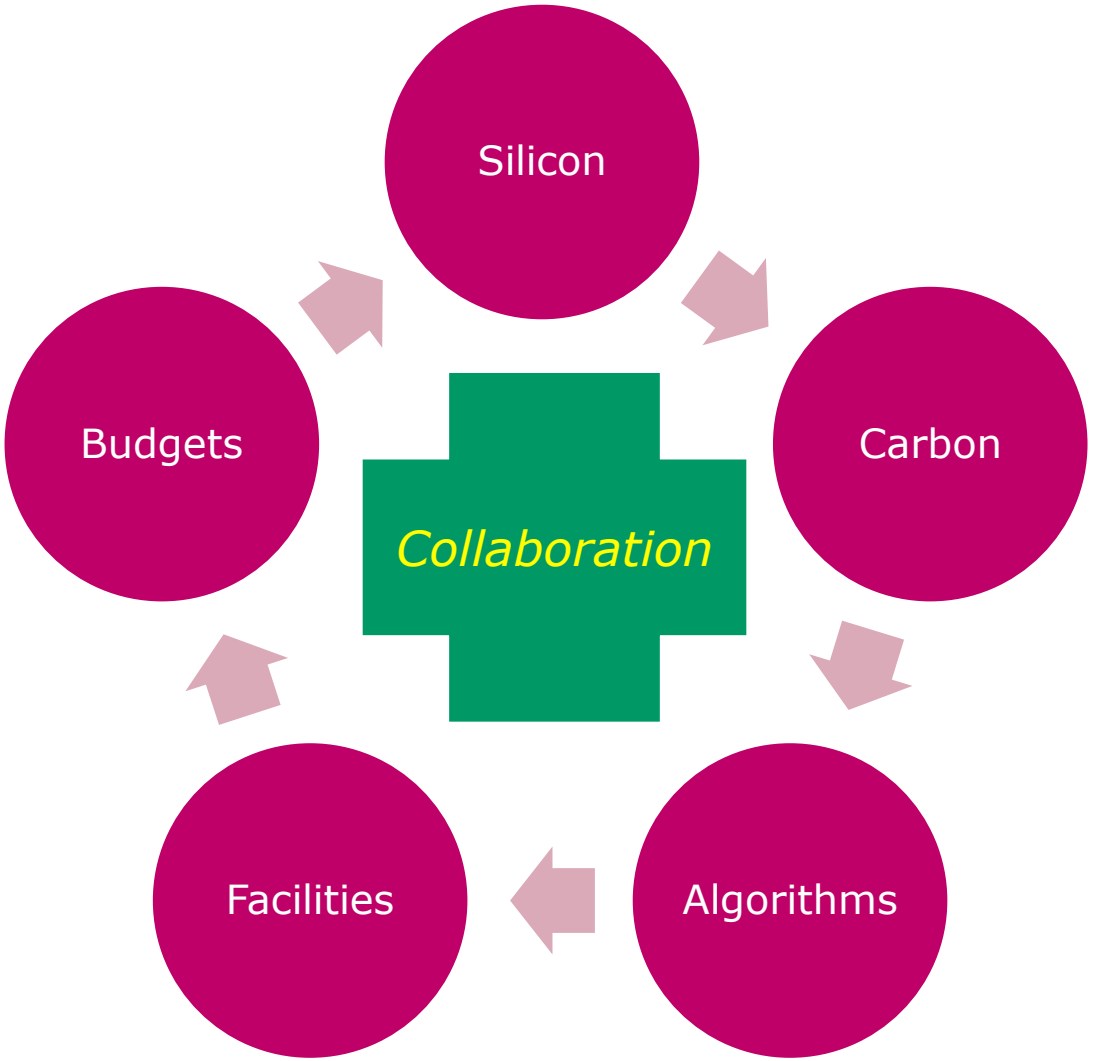

AMD Collaboration Resources

AMD
Smarter Choice

- Green Grid

  http://www.thegreengrid.org/home

- Developer Pages

  http://developer.amd.com/

- Torrenza Forum (accelerators)

  http://enterprise.amd.com/us-en/AMD-Business/Technology-Home/Torrenza.aspx

- Lightweight Profiling for increased parallellism

  http://developer.amd.com/lwp.jsp

- HyperTransport interconnect

  http://www.hypertransport.org/

My Prediction…

AMD
Smarter Choice

Texas A&M 38

OSU 13

Gig 'Em Aggies!

Disclaimer and Trademark Attribution

DISCLAIMER

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 OR FOR THE PERFORMANCE OR OPERATION OF ANY PERSON, INCLUDING, WITHOUT LIMITATION, ANY LOST PROFITS, BUSINESS INTERRUPTION, DAMAGE TO OR DESTRUCTION OF PROPERTY, OR LOSS OF PROGRAMS OR OTHER DATA, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.