AMD
Smarter Choice

Future Directions in Processor Technology

David Rich
Director, Commercial Solutions – AMD
President, HyperTransport Consortium

- excerpt -

Agenda

Macro Trends

Now (Quad Core)

The Future

- Bandwidth, Latency, Directions…

Trends / What can you do with CMOS?

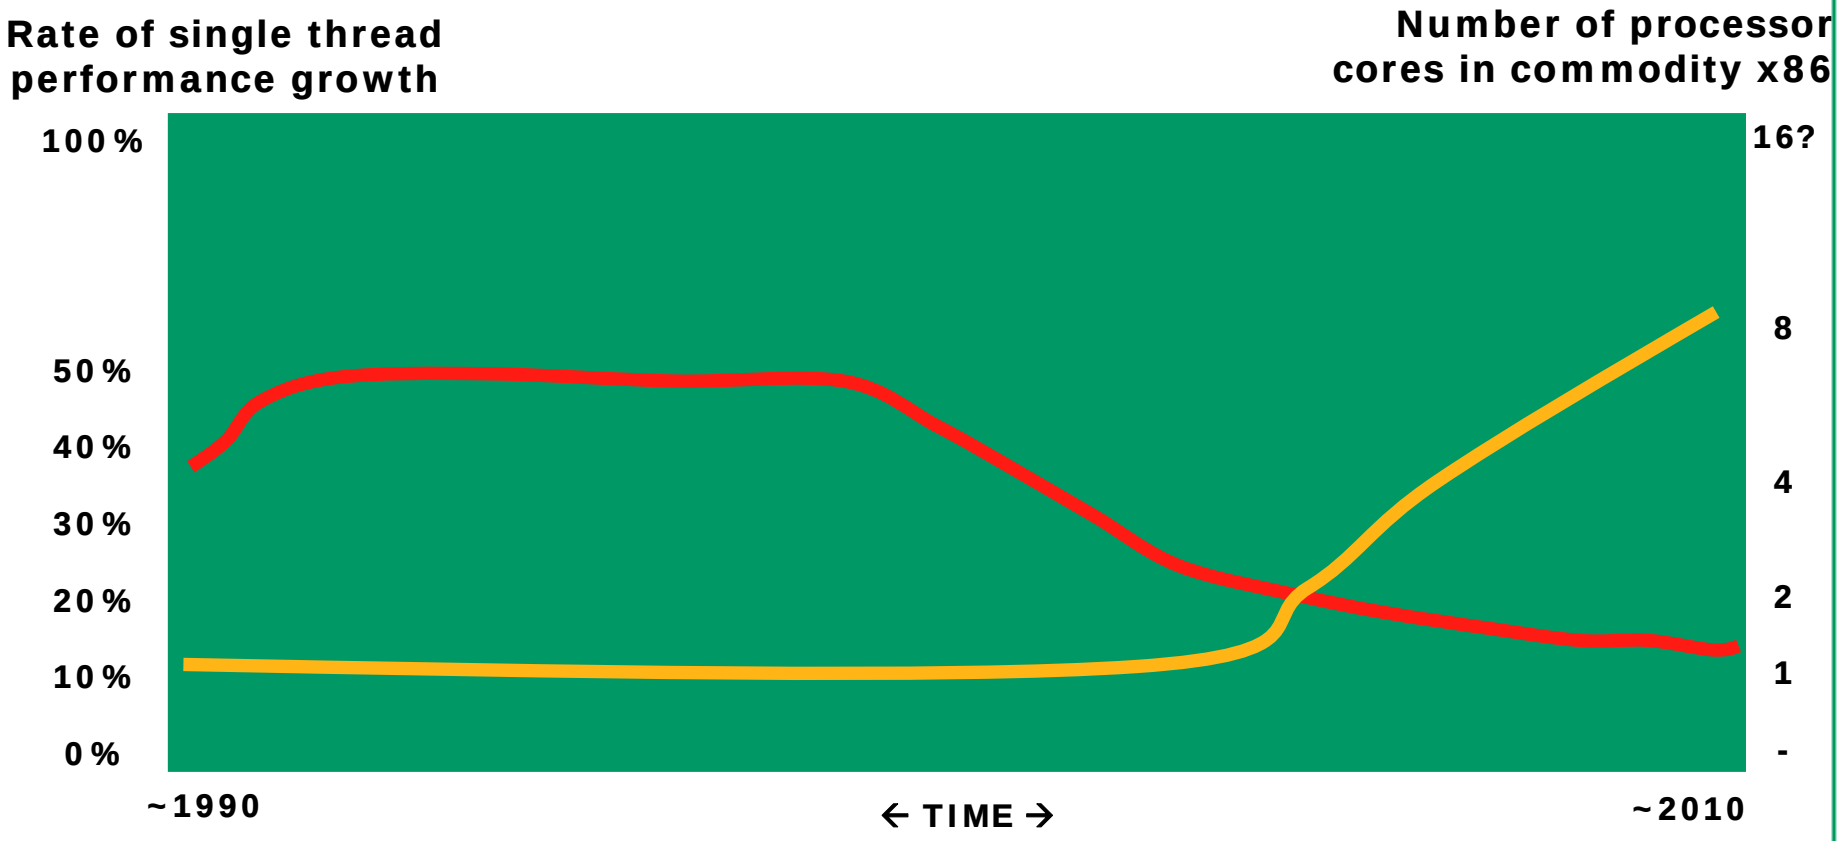

Graph is illustrative only

Trends / What can you do with CMOS?

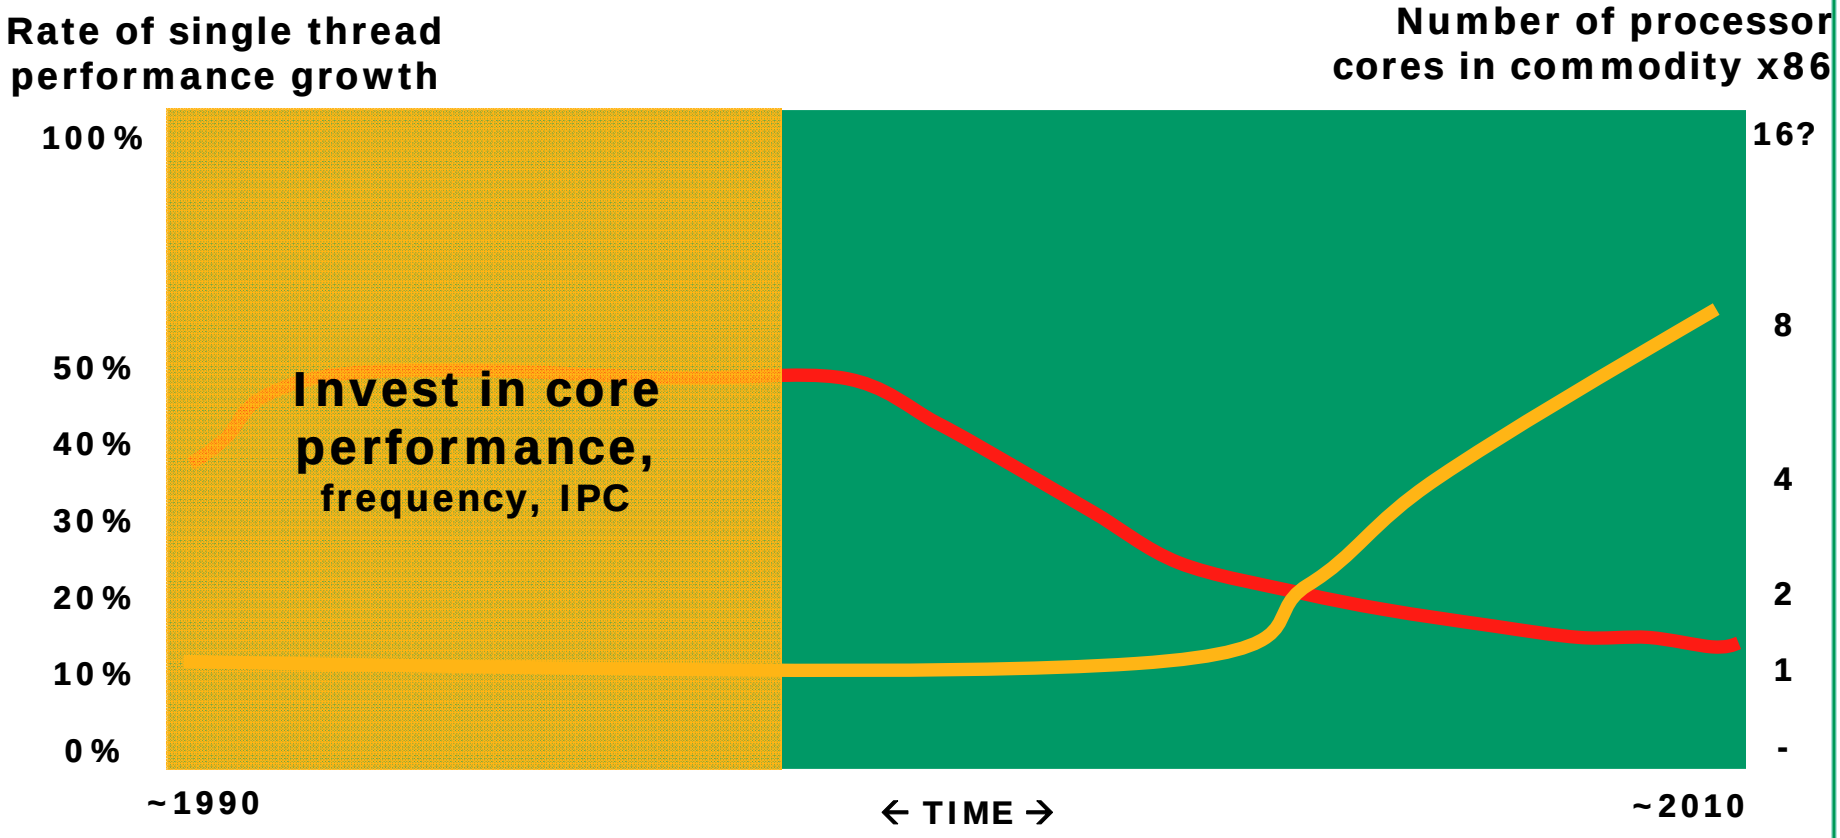

Graph is illustrative only

Trends / What can you do with CMOS?

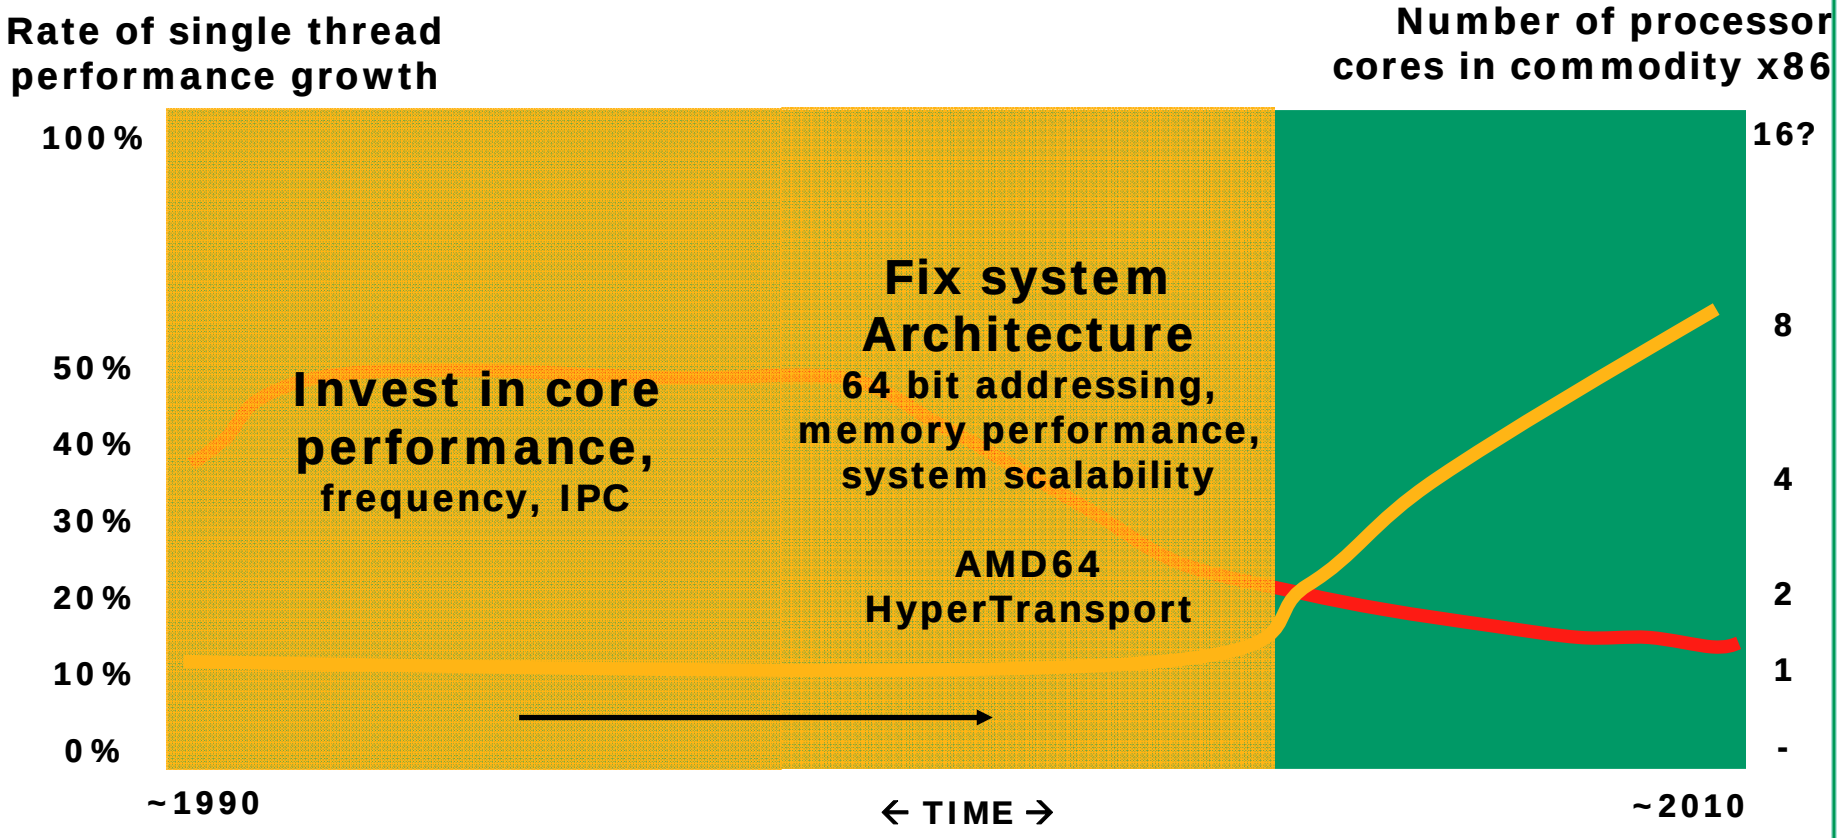

Graph is illustrative only

Trends / What can you do with CMOS?

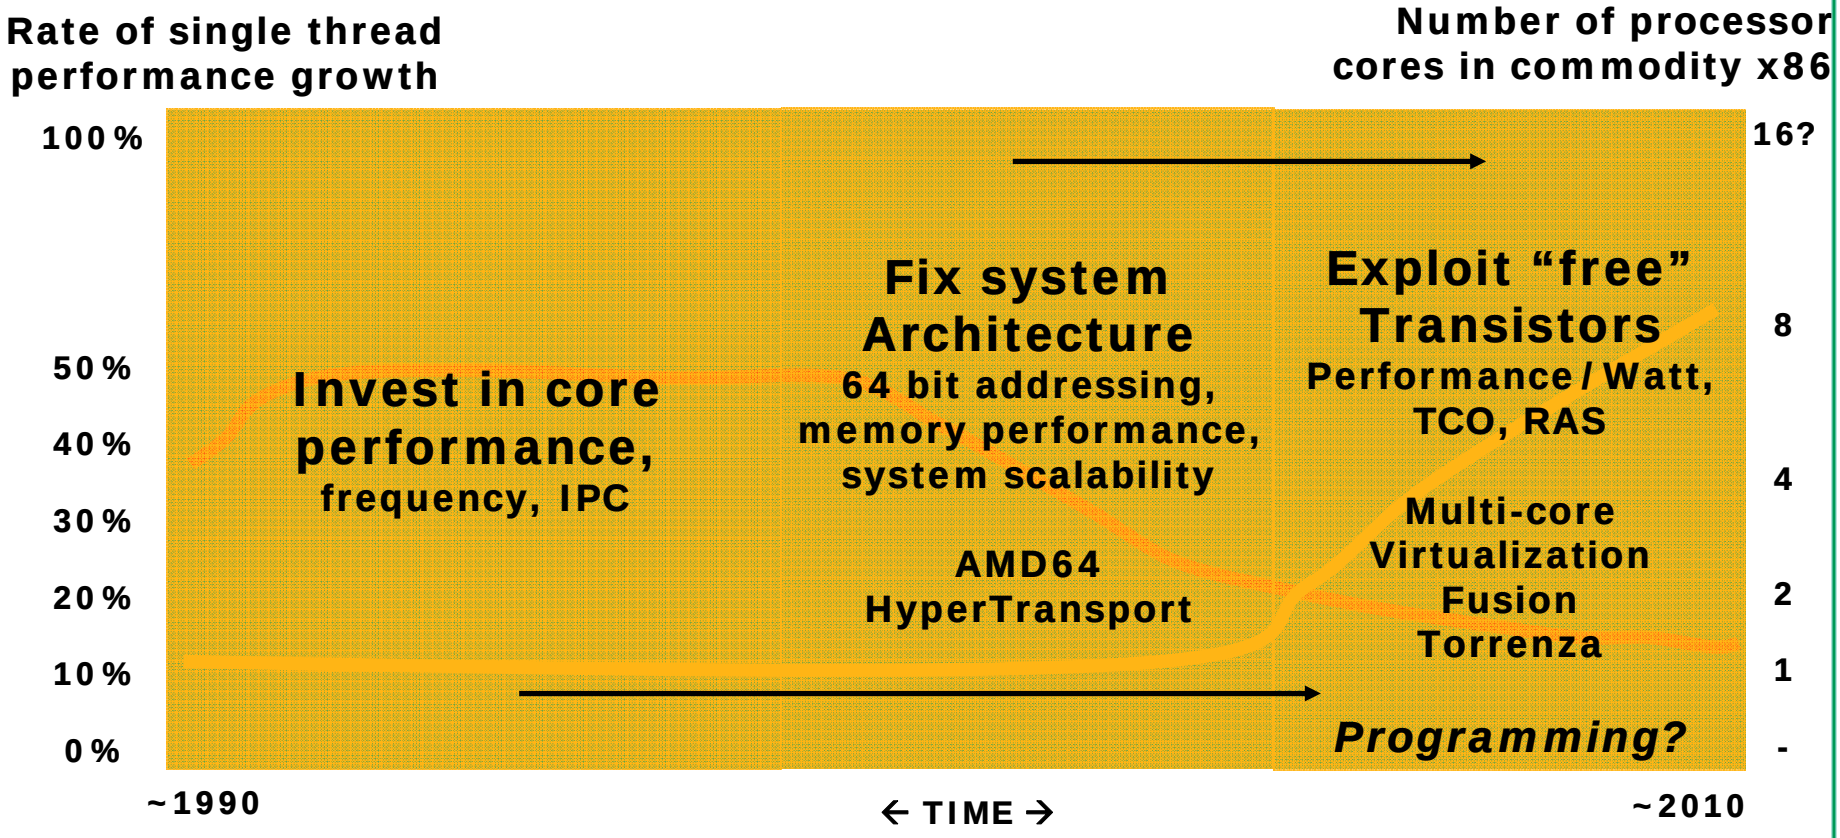

Graph is illustrative only

Single → Dual → Quad, the AMD way

AMD Opteron™ Processor Platform Longevity

AMD Smarter Choice

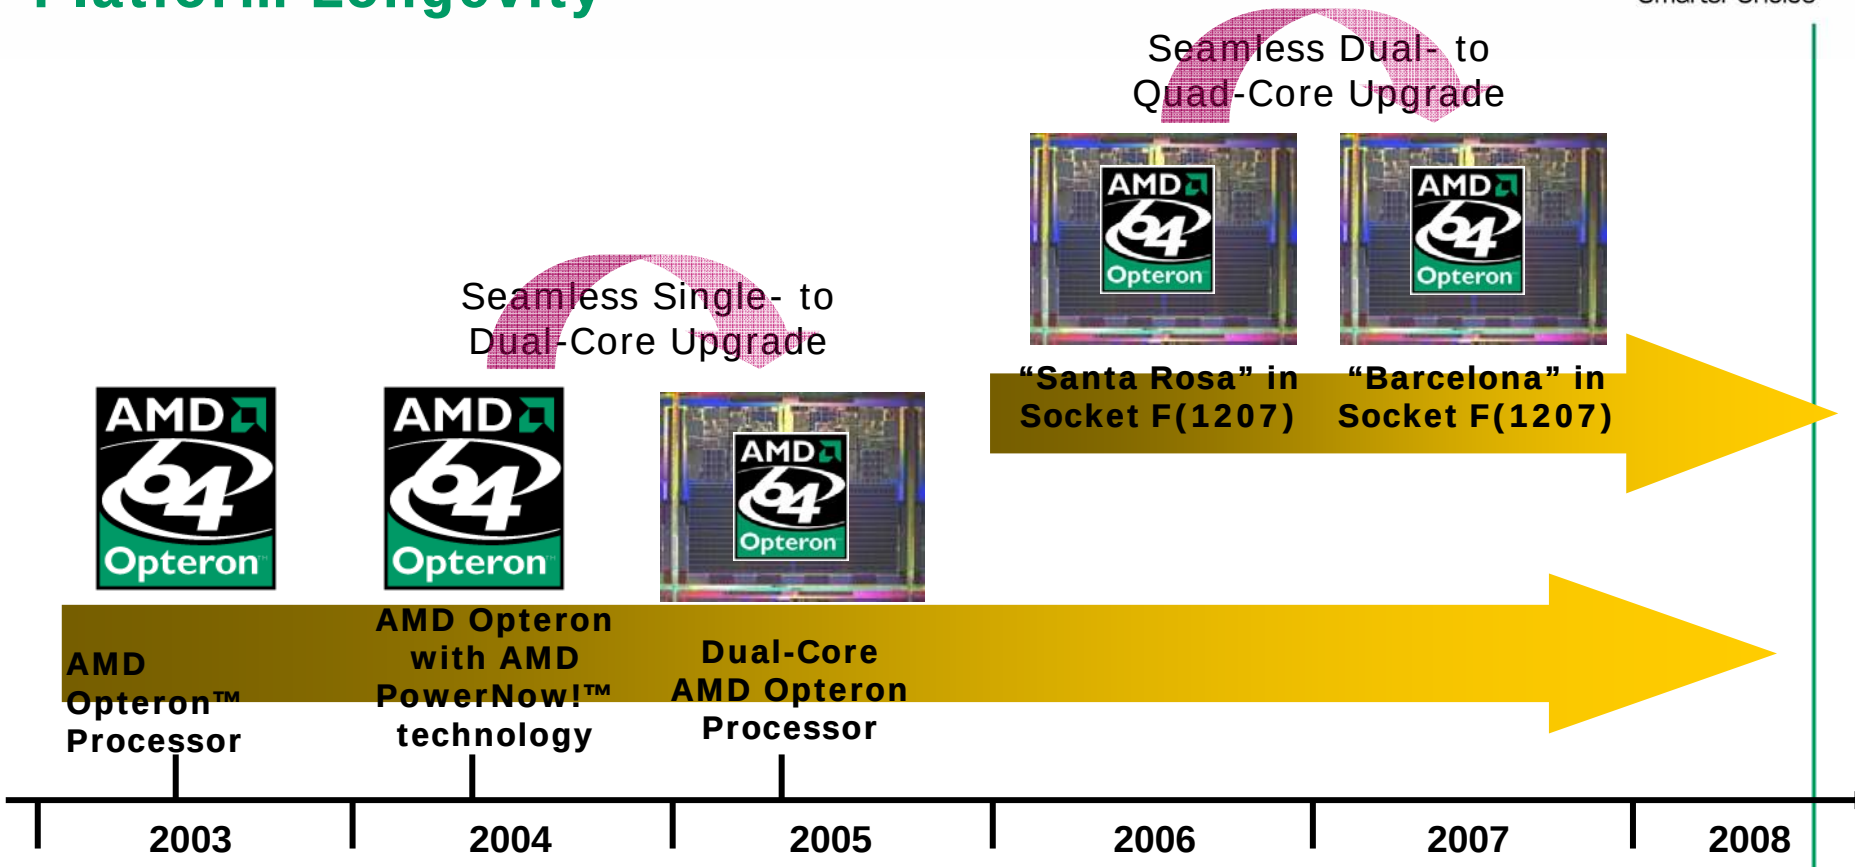

Socket 940

- Introduced in April 2003
- Continued support into 2008

Socket F(1207)

- Introduced in 2006
- Continued support through 2009

6+ years of AMD Opteron... 2 sockets

AMD Barcelona quad-core Processor

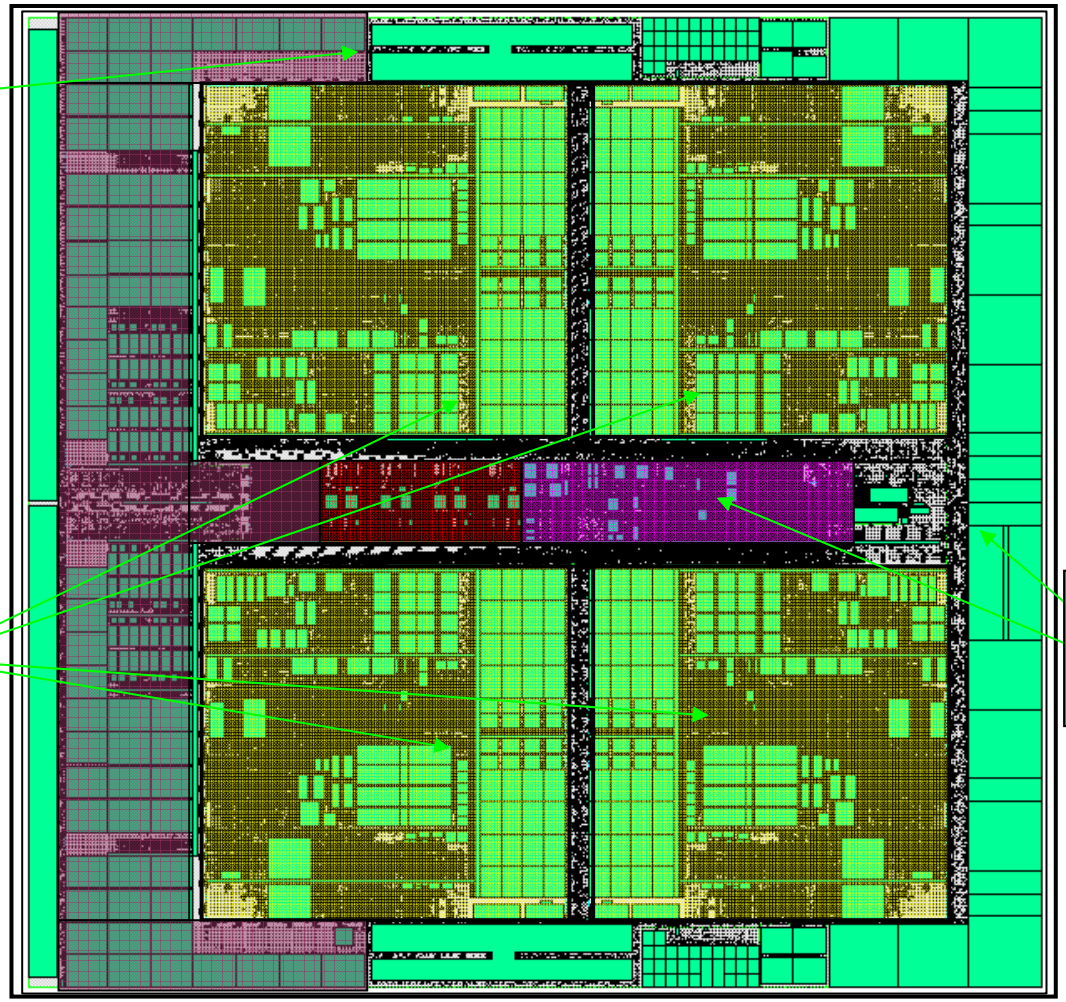

Comprehensive Enhancements for SSE128

Double Flops

IPC-enhanced CPU cores

TLB and latency

Memory controller

DRAM Bandwidth

Quad-Core AMD Opteron™ Processors

AMD Smarter Choice

More than just four cores

- **Significant CPU Core Enhancements**
- **Significant Cache Enhancements**

World-class performance

- **Native Quad Core**
  - Faster data sharing between cores
- **Enhanced AMD-V™**
  - Nested paging acceleration for virtual environments

Reducing total cost of ownership

- **Performance / Watt** leadership
  - Consistent 95W thermal design point
  - Low power 68W solutions
- **Drop-in upgrade**
  - Socket F compatibility – BIOS upgrade
  - Leverage existing platform infrastructure
- **Common Core Architecture**
  - One core technology top-to-bottom
  - Top-to-bottom platform feature consistency

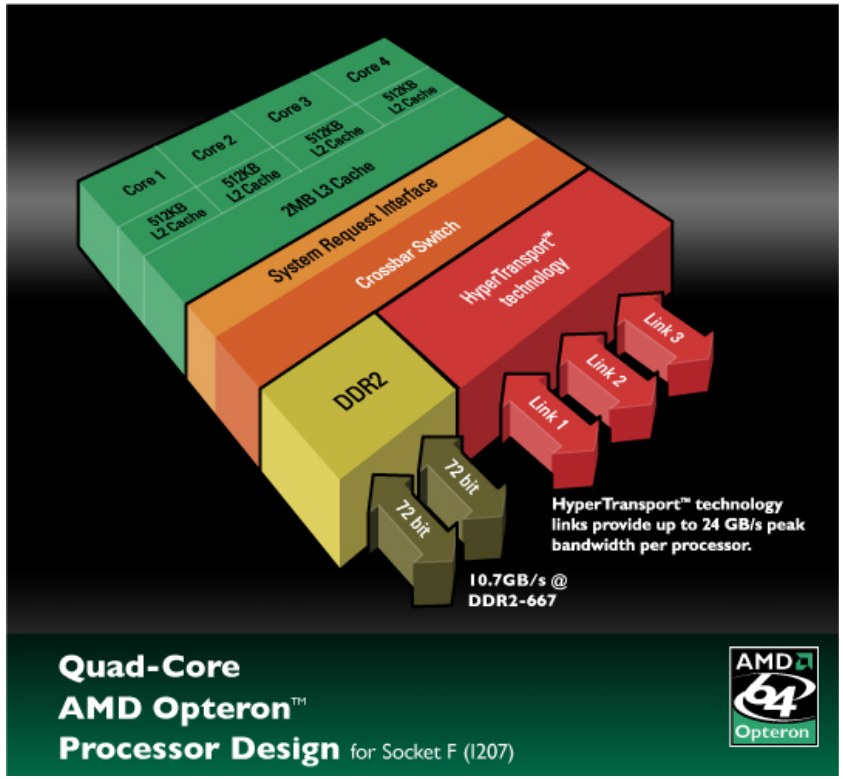

System Bandwidth Also Important…

AMD Smarter Choice

HyperTransport™ CONSORTIUM

State-of-the-Art Specifications

HyperTransport 3.0
State-of-the-Art Specifications

Same Features as HT 2.0 Plus:

- 1.8 GHz, 2.0 GHz, 2.4 GHz and 2.6 GHz Clock Support
  - 41.6 GB/s Aggregate Bandwidth
  - 20.8 GB/s (166.4 Gb/s) per Link
- DC Operating Mode Enhancements
- AC Operating Mode (Optional)
  - Supports Applications Requiring Greater Signal Interconnect Distance
    - Cables
    - Backplanes
    - Larger Physical Systems
    - Chassis-to-Chassis Connections
- DC/AC Auto-Configuration
- Link-Splitting/Un-Ganging Mode (Optional)
  - Auto-configuration of Bi-Mode 2x8 or 1x16 Links
- Hot Plugging
  - Backplanes Applications
- Power Management Enhancements (Optional)
  - Support Dynamic Link Frequency and Width
- 100% Backward Compatibility
  - Auto-Configuration at Boot-Up with Minimum Spec Common Denominator Selection

Double speed

Form factor flexibility

Lower latency

www.hyptransport.org for more information

SW enablement of HW technologies …

AMD Smarter Choice

Hardware	**Software**
Single CPU Core	Optimizing compilers (Compilers)
Tightly-coupled SMP	Locks & Synchronization (OS, **Apps**) Sophisticated VMM (OS)
Collection of SMP nodes (NUMA, NUMBW exposed)	Process & memory affinity (OS) Virtualization (Hypervisor)
Accelerators & co-processors	APIs, standards, libraries (**Apps**)
Heterogeneous SMP	Capability aware scheduling (OS)
Clustering	MPI, Cache Fusion, Parallel DB (**Apps**)
Grid	Fault tolerant “meta apps” (**Apps**)
Extreme NUMA (due to chip stacking technology)	Algorithm-level cache blocking (**Apps**)? Exposed ping-pong buffering (**Apps**)?

Perhaps capacity may be large enough so many applications work with no changes?

A Diverse Future Is Possible (Likely?)

AMD Smarter Choice

AMD's **Torrenza** Program supports an ecosystem of accelerators and co-processors

- Public supporters →

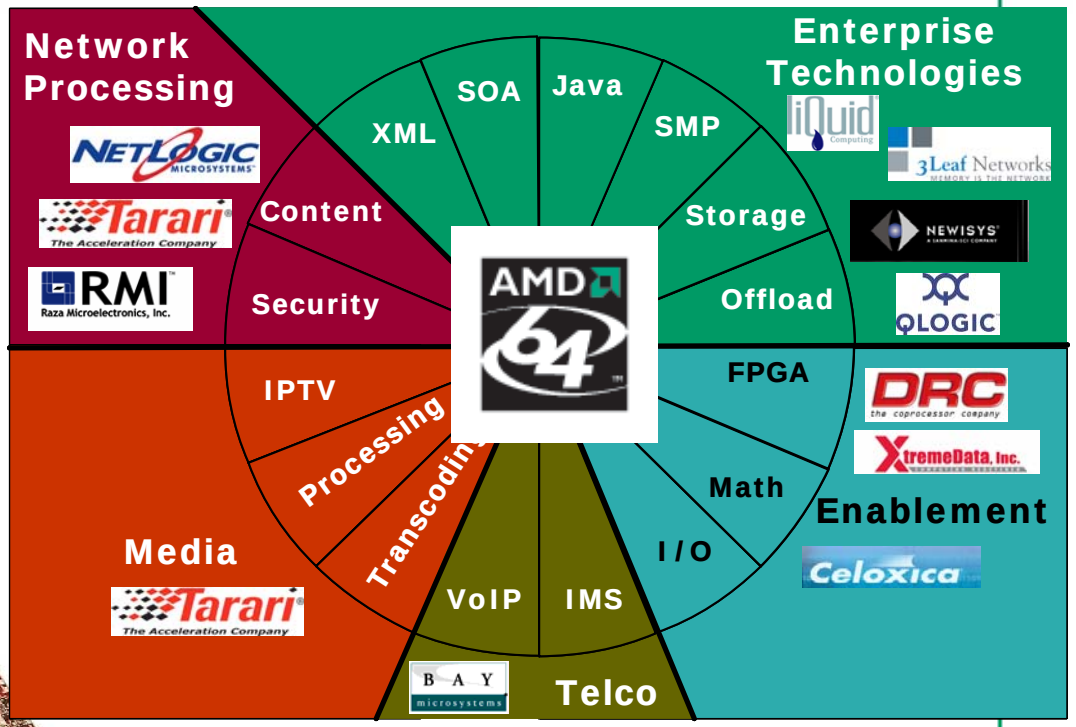

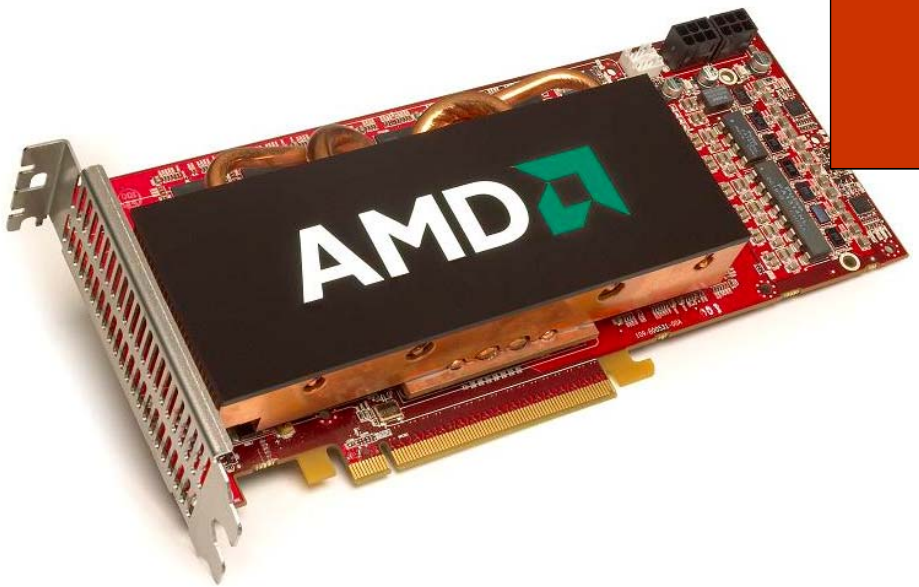

AMD's **Stream Processors** exploit GPU technology for high performance computation

Continuum of Solutions

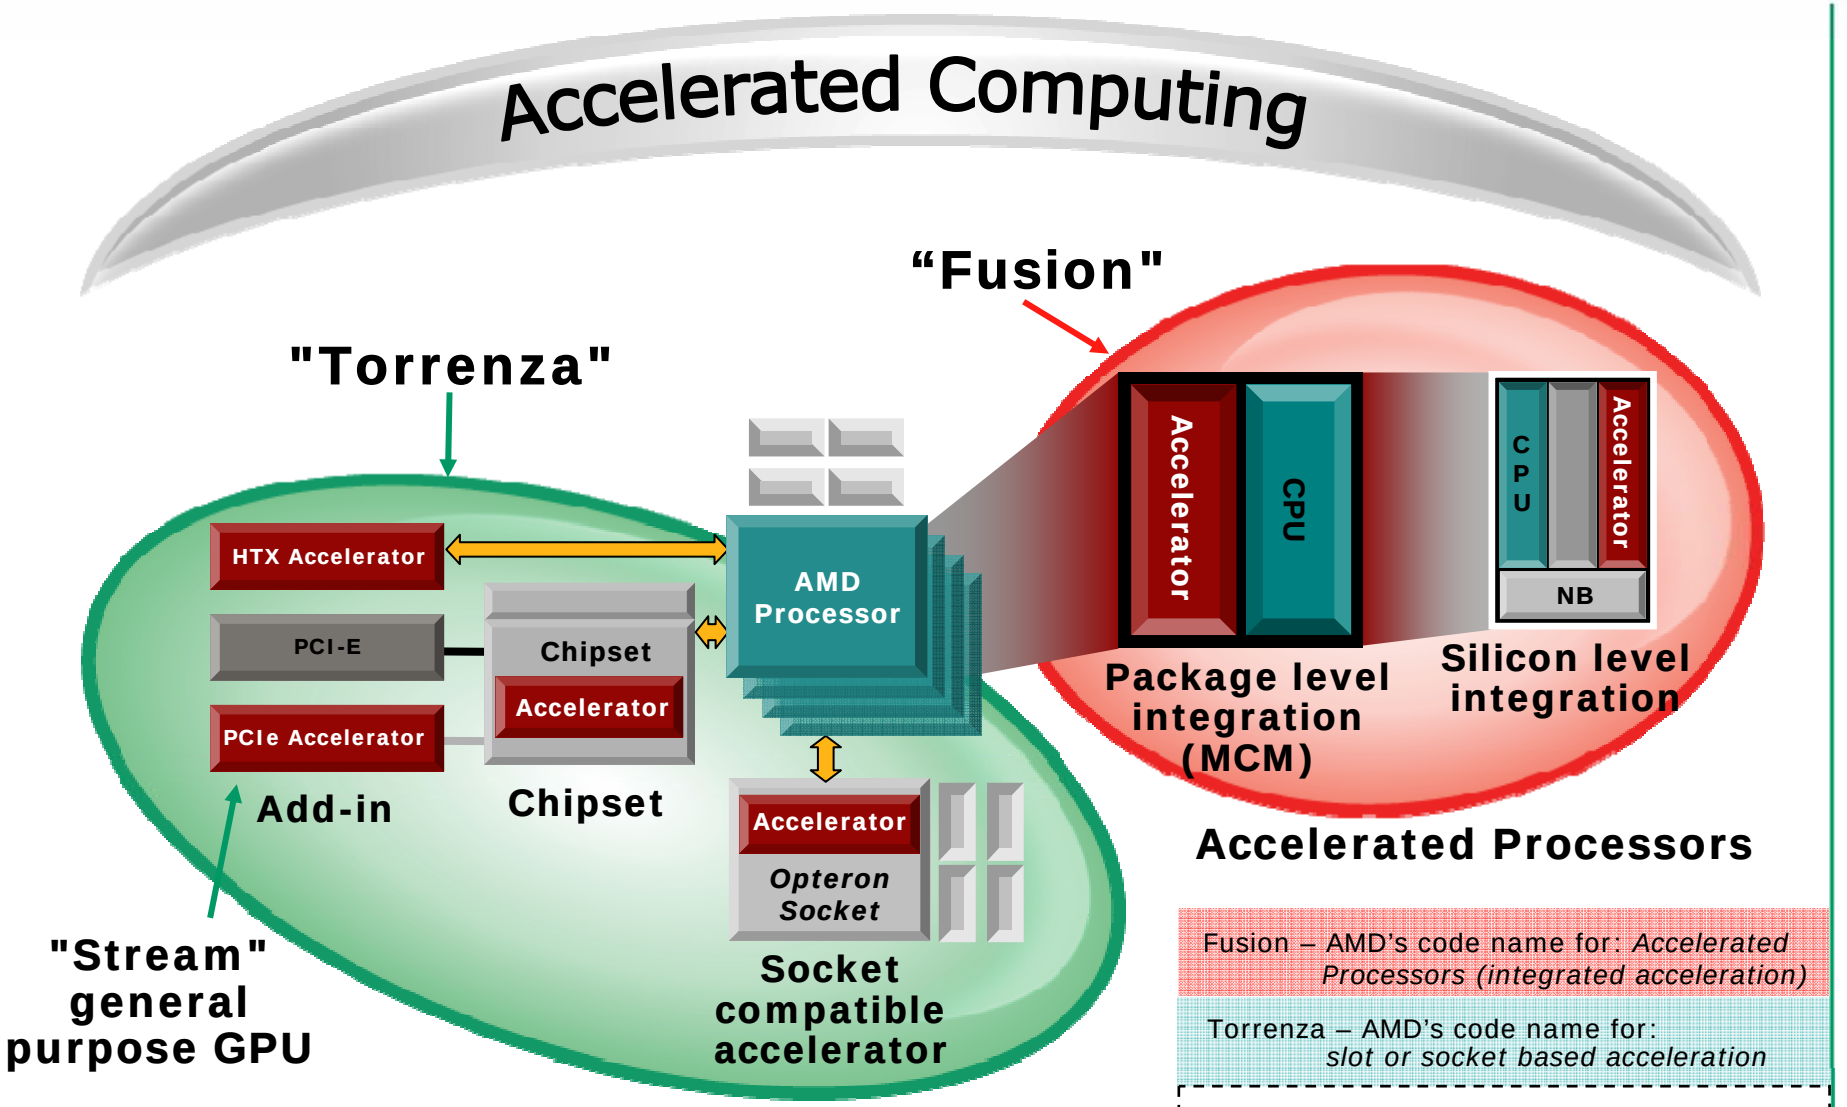

Fusion – AMD's code name for: *Accelerated Processors (integrated acceleration)*

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

Stream – *Specific example of a GPGPU accelerator under Torrenza*

Summary

AMD's technology deployment decisions and product definitions are the result of a customer-centric process

- Right technology at right time
- Solutions solve customer problems and offer real value
- Economic considerations play an important role
  - ASPs tend not to rise! Can't bury cost into system-level…

Key Challenges for future "extreme capability" systems:

- Programmability
- Memory: Tricky balance between BW, Latency, Capacity, Power & Cost
- Balancing data center level performance/watt attributes

Infrastructure Information

We encourage our customers & partners to visit AMD's public website for infrastructure items including:

- System Component Information
  - *Recommended Motherboards*
  - *Thermal Solution Guidelines*
  - *Thermally Tested Server Solutions*
  - *Memory Guidelines*
  - *Power Supply Guidelines*
  - *Tower Chassis Guidelines*

  http://www.amd.com/configuration

- Open Platform Management Architecture (OPMA) specification
  www.amd.com/opma

- Online Processor Quick Reference Guide
  http://www.amd.com/processorquickrefguide

- AMD64 Ecosystem
  www.amd.com/amd64ecosystem

- AMD Lead-Free Initiative
  http://www.amd.com/leadfree

- Performance
  - www.amd.com/opteronperformance

- Pricing
  http://www.amd.com/pricing

- Model Numbers
  http://www.amd.com/modelnumbers

- Processor Diagram
  http://www.amd.com/architecutral-features

- Where to Buy
  http://shop.amd.com/us-en/platforms.aspx

Thank You!

David Rich
David.Rich@AMD.com

Trademark Attribution