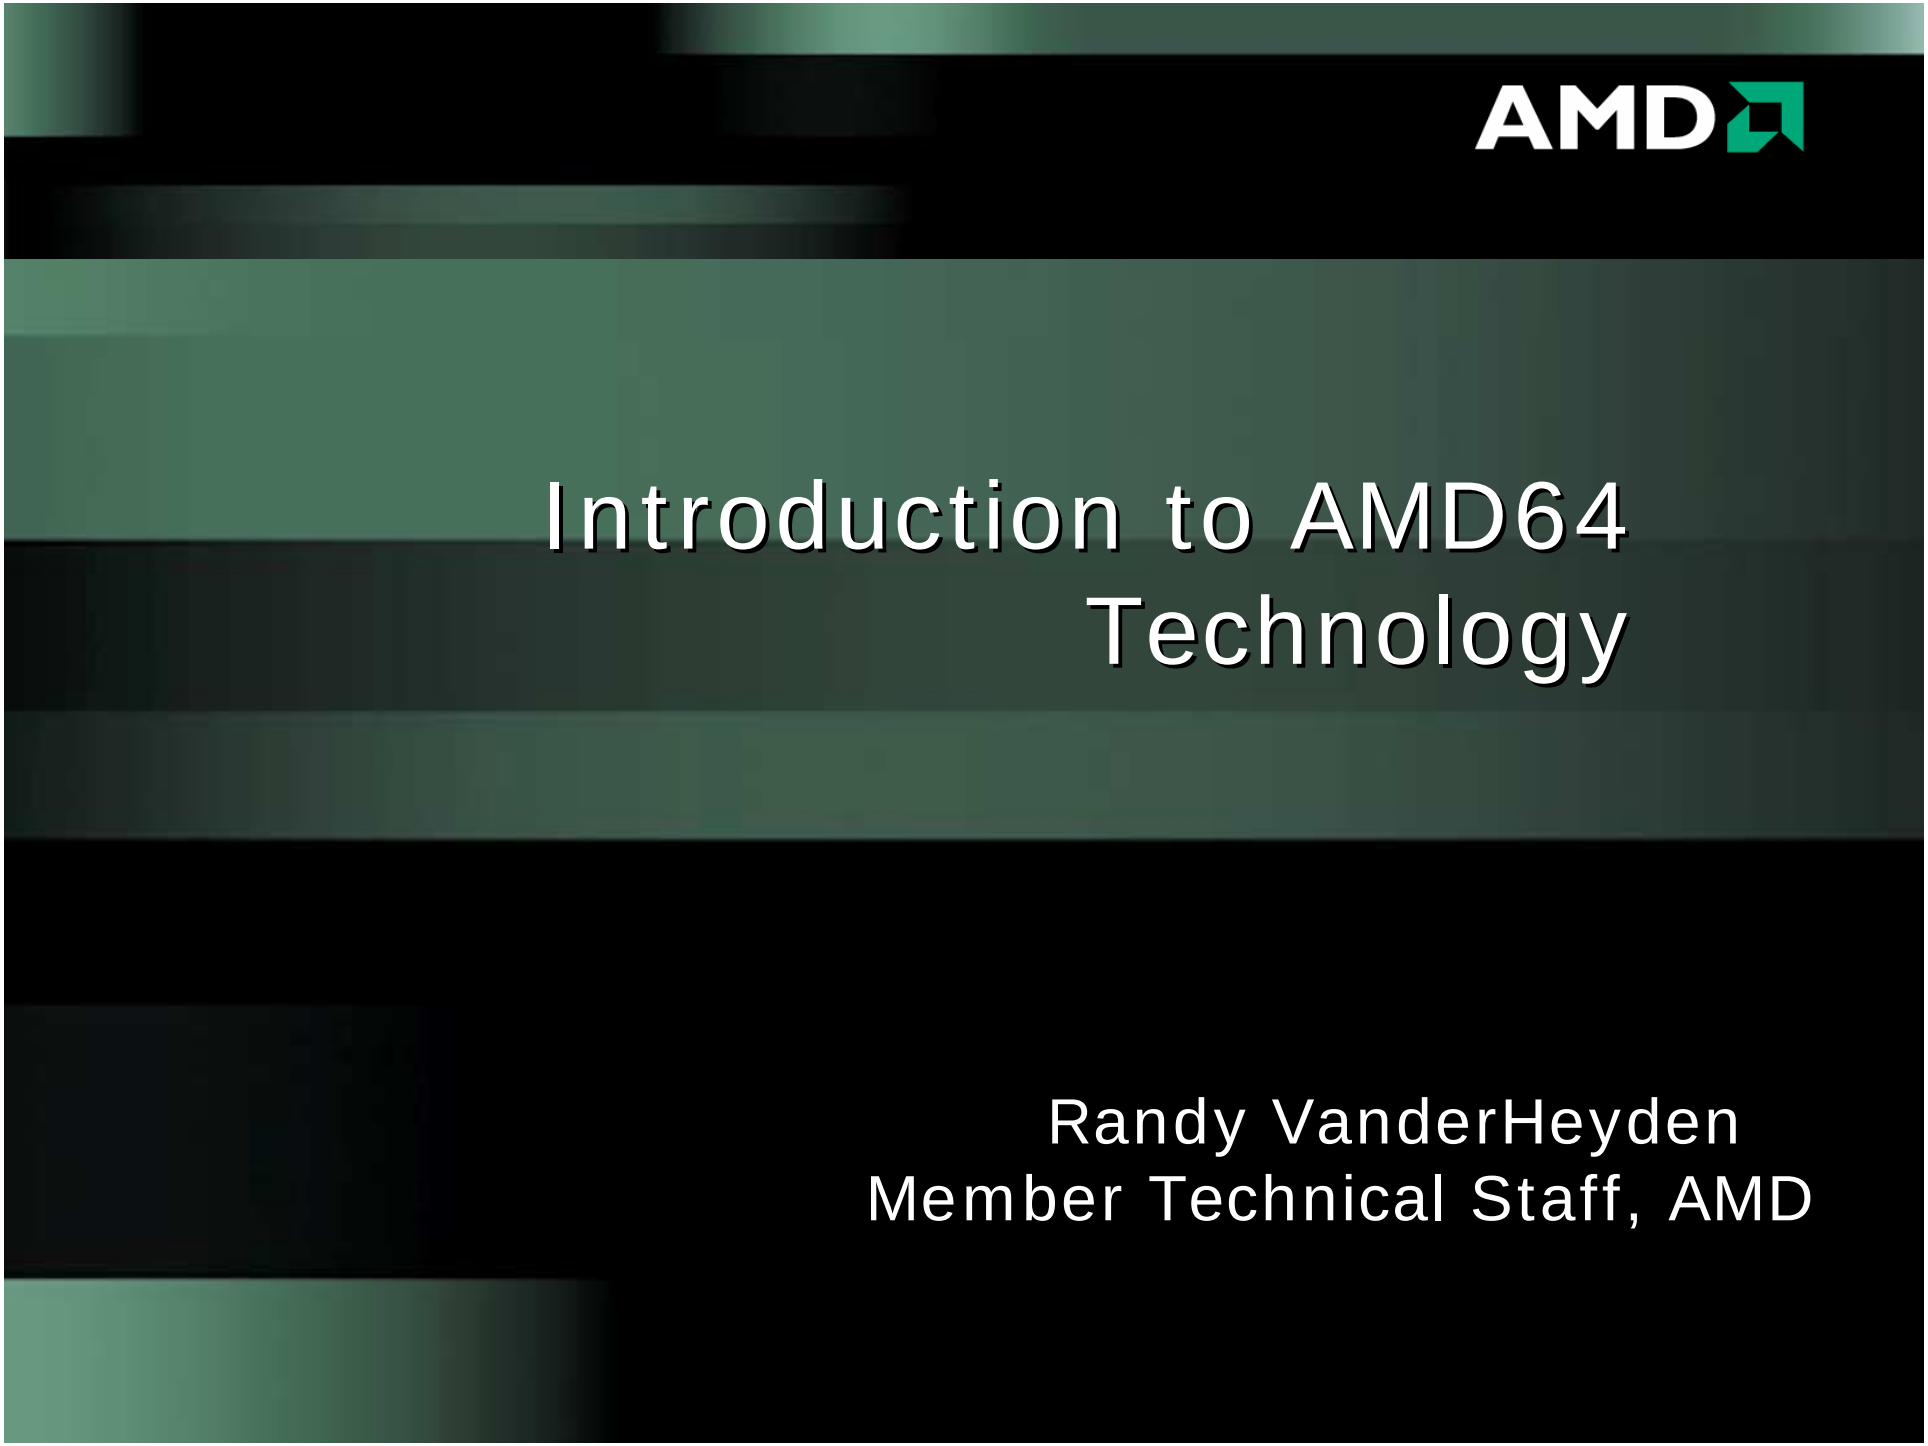
AMD
Introduction to AMD64 Technology
Randy VanderHeyden
Member Technical Staff, AMD

AMD Overview

A leading global supplier of innovative semiconductor solutions for the personal and enterprise computing, communications and consumer electronic markets

Founded: 1969

Headquarters: Sunnyvale, California

Employees: 15,000 worldwide

Sales Mix: 80% international

2004 Revenue: $5 billion

2004 net income: $91 million

2004 Q4 Revenue: $1.26 billion

2004 Q4 op income: $20 million

- Microprocessors: $730 million
- Flash memory: $504 million

AMD Worldwide Operations

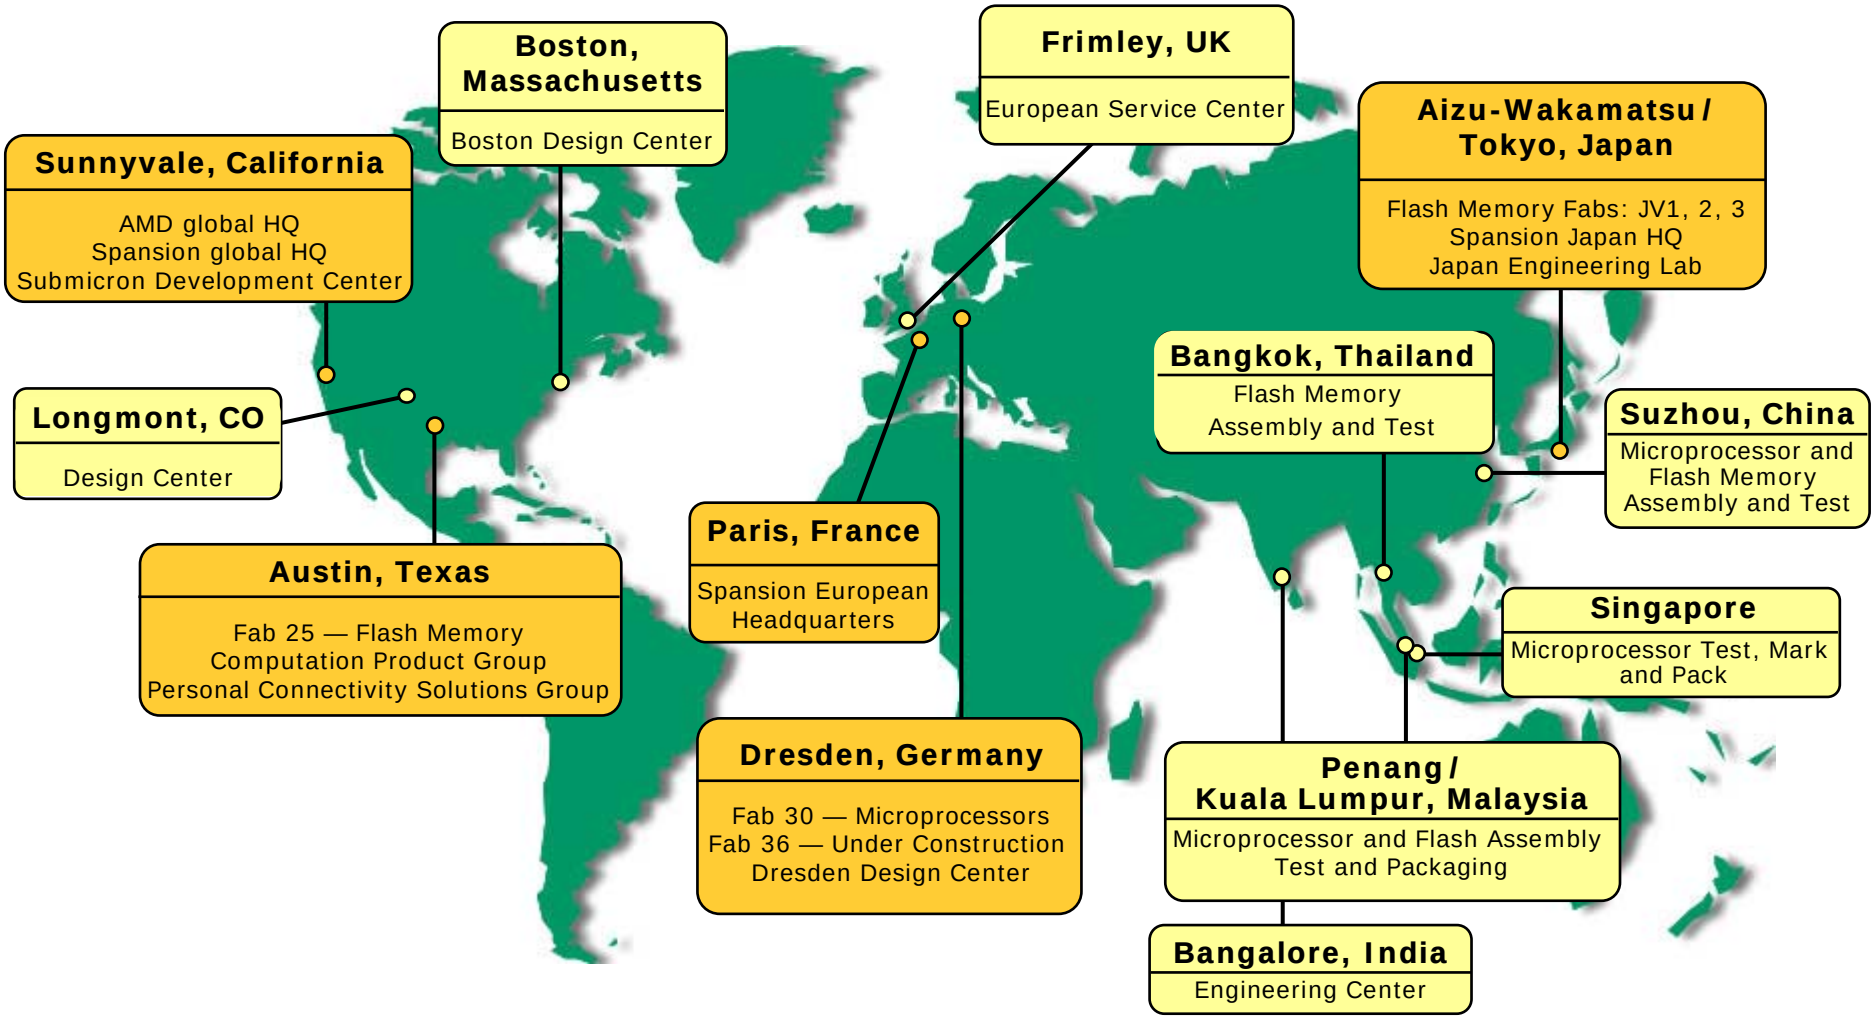

AMD Opteron™ Processor Tier 1 Server Performance Wins

Sun microsystems

- #1 4P SPECweb®99_SSL
- #1 4P SAP SD 2-Tier (x86 Linux)
- #1 4P SPECompM2001

Sun Fire V40z

- #1 2P SAP SD 2-Tier (x86 Linux)
- #1 Dual Node SPECjAppServer®2002
- #1 Multi Node SPECjAppServer®2002

Sun Fire V20z

IBM

- #1 2P TPC-H™ 100GB clustered Performance
- #1 2P TPC-H 300GB clustered Performance

IBM eServer 325

hp invent

- #1 4P TPC-C™ $/tpmC
- #1 4P SAP SD 2-Tier (Windows®)
- #1 4P SPECint®_rate2000 (Linux)
- #1 4P SPECfp_rate2000 (x86 Linux)
- #1 4P TPC-H™ 1000GB clustered performance
- #1 4P Lotus NotesBench R6.0

HP ProLiant DL585

HP ProLiant DL385

- #1 2P/2U TPC-C
- #1 2P/2U SAP SD 2-Tier (Windows)
- #1 1P SPECint®_peak2000 (Linux)

HP ProLiant BL25p

- #1 2P Blade SAP SD 2-Tier (Windows)
- #1 1P SPECfp®_peak2000 (x86 Linux)
- #1 2P SPECint®_rate2000 (Linux)
- #1 2P SPECfp®_rate2000 (x86 Linux)

AMD Opteron™ Processor Tier 1 Workstation Performance Wins

AMD Opteron™ processor-based workstations own the #1 performance score on every industry-standard SPEC graphics/applications benchmark

- #1 1P SPECviewperf 8.0
- #1 1P SPECapc(SM) for 3ds max 6
- #1 1P SPECapc(SM) for Maya 6

HP xw9300

- #1 1P SPECapc(SM) for Maya 5

IBM IntelliStation A Pro

- #1 2P Ensight
- #1 1P SPECompM®2001 (x86)

Sun Java Workstation W1100z and W2100z

AMD Opteron™ Technology and Features

AMD64 Technology

- Introduced x86 64-bit extension technology to the world

Integrated Memory Controller

- No longer in the Northbridge, the memory controller is now in the processor
- Memory controller and queues run at CPU core speeds, reducing latencies by over 30%

HyperTransport™ Interface

- High-bandwidth, low-latency pt-to-pt coherent interconnect

Scalability

- Glueless architecture to scale from one to four processors
- Resources (memory, I/O) scale with number of CPUs

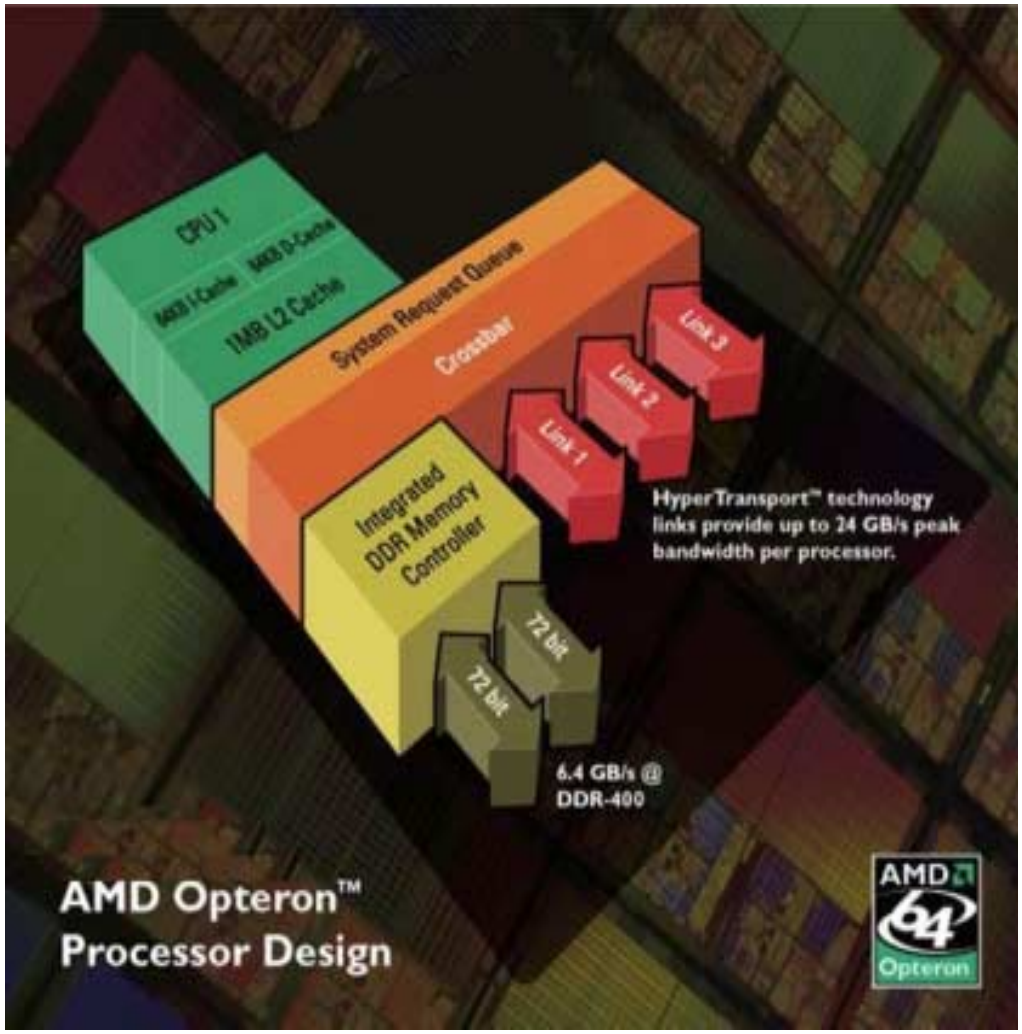

Integrated Memory Controller

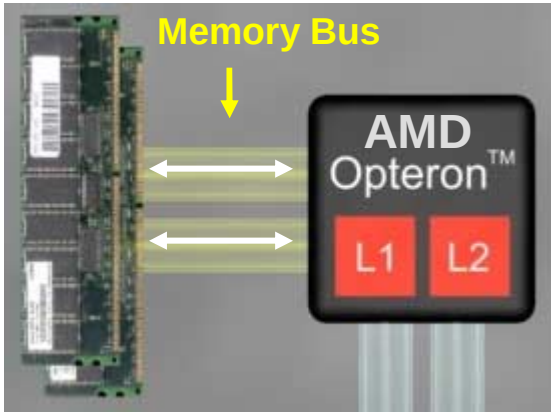

Integrated Memory Controller

- Runs at the CPU core frequency
- Becomes more efficient as the CPU frequency increases
- Each CPU has dedicated bandwidth for memory access
- Bandwidth increases with addition of CPU's

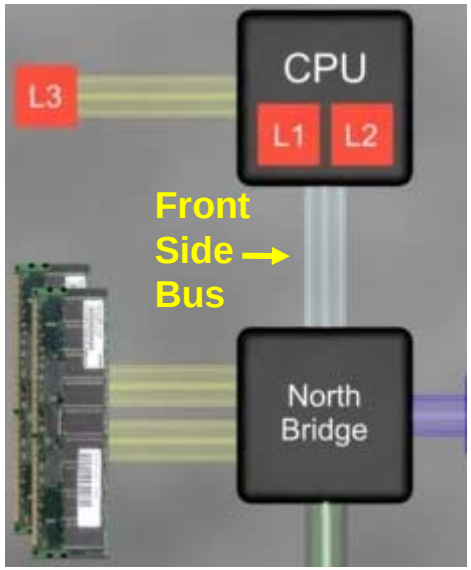

External Memory Controller

- Memory performance scales with the FSB frequency
- Additional CPU's must share memory bandwidth

HyperTransport™ Technology

- HyperTransport™ technology is
  - High-speed, low pin-count, asynchronous, chip-to-chip board level interconnect
  - Proven, industry-standard technology in production today
- HyperTransport is *not*
  - A replacement for PCI or its roadmap (PCI-X, PCI-X/DDR, PCI-3.0)
  - A networking fabric
- For AMD Opteron processors there are three HyperTransport links per processor

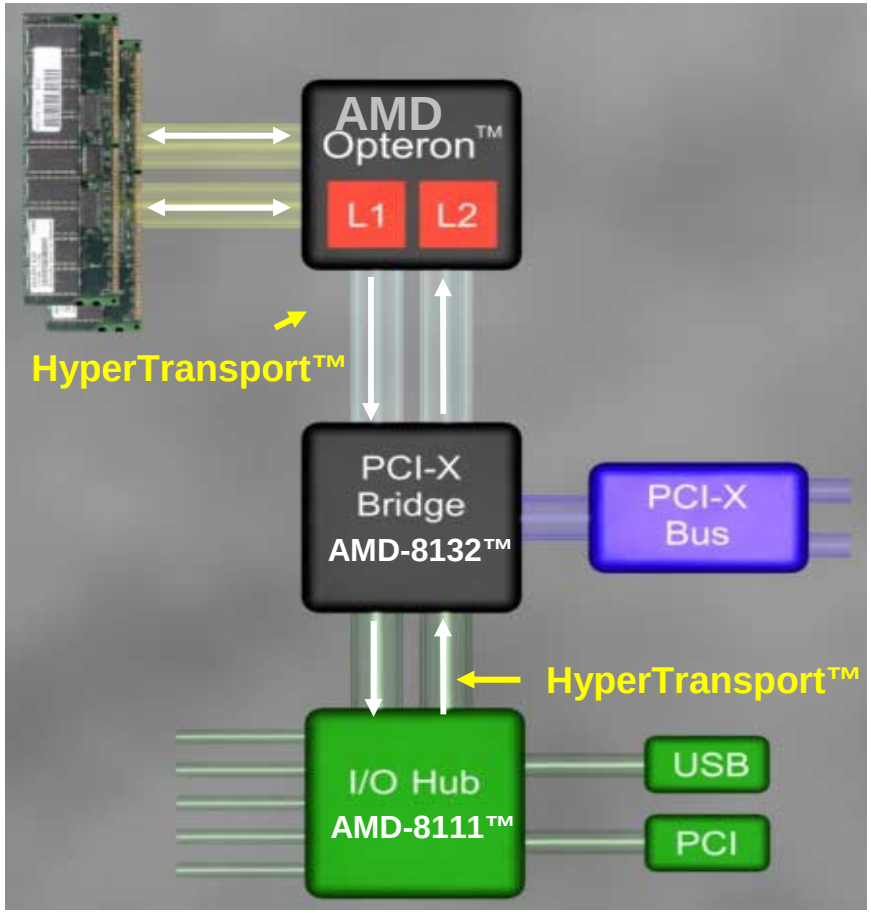

HyperTransport™ Technology

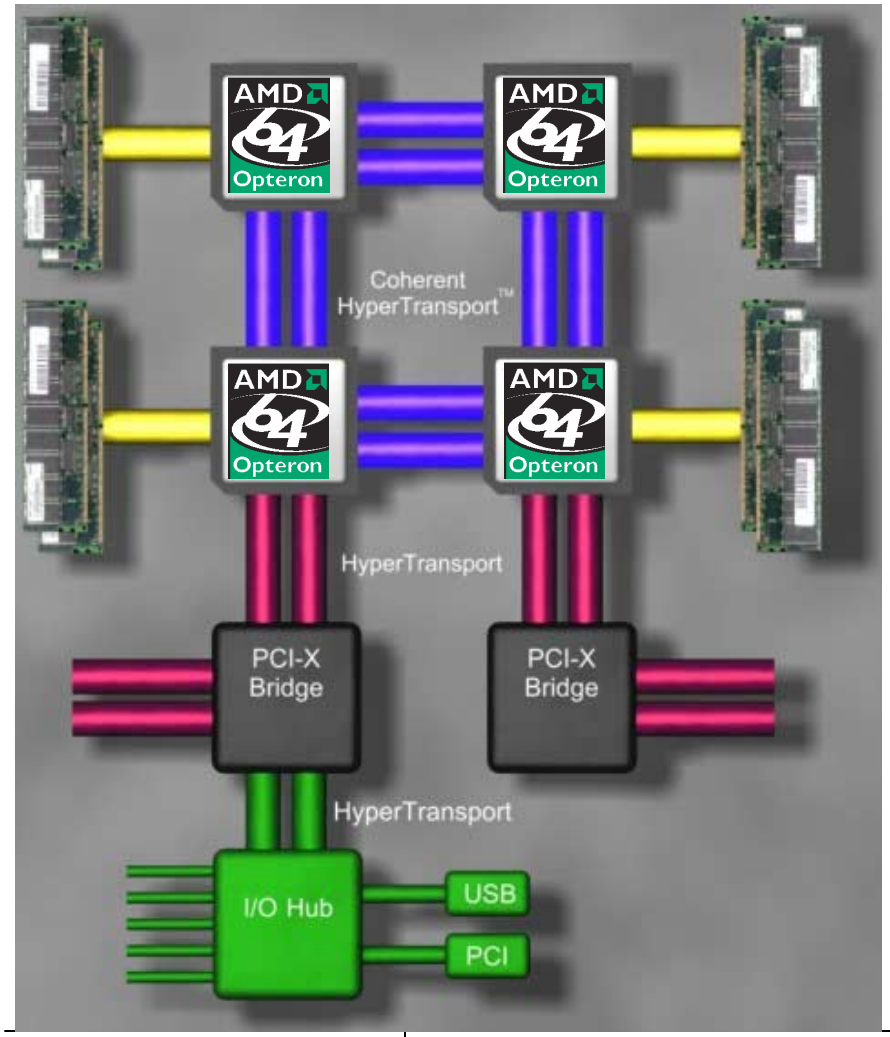

"Glueless" Multi-CPU Design

- Point-to-point link
- Allows for "glueless" multi-CPU designs
- Each CPU adds memory and HyperTransport bandwidth
- Three (coherent capable) HyperTransport links per AMD Opteron™ CPU (8xx Series)
- Scales to an 8-way without additional glue logic

AMD

Non-Uniform Memory Access (NUMA)
Memory is distributed across a number of CPUs

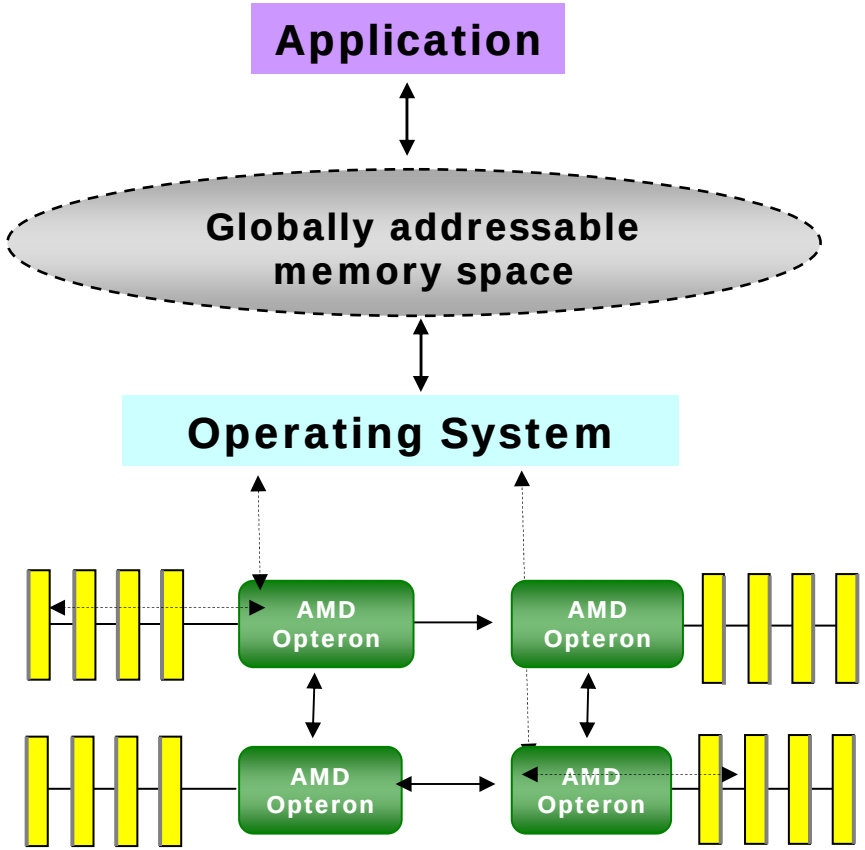

Software accesses data belonging to any processor via the global address space

Memory is initialized into globally addressable physical memory space with processors maintaining cache coherency across this space

NUMA-aware OS assigns threads from same process to the same NUMA node

Each processor has local memory and uses HyperTransport for high-speed access to non-local memory

More NUMA

- 64-bit Windows® implements ccNUMA support
  - processes and memory are allocated intelligently
  - apps: allocate per-thread memory in the thread
  - use ccNUMA API funcs if you need more control over thread or process assignment:

    GetNumaHighestNodeNumber, SetThreadAffinityMask
- Performance benefits on all 64-bit Windows OS
  - XP Pro 64-bit, Server 2003 64-bit, Advanced Server 2003 64-bit (and also on 32-bit Advanced Server 2003)
- One configuration tip! If “Bank Interleaving” is an option in the BIOS, disable it for ccNUMA aware OS’s.

MSDN Has Useful Information On ccNUMA

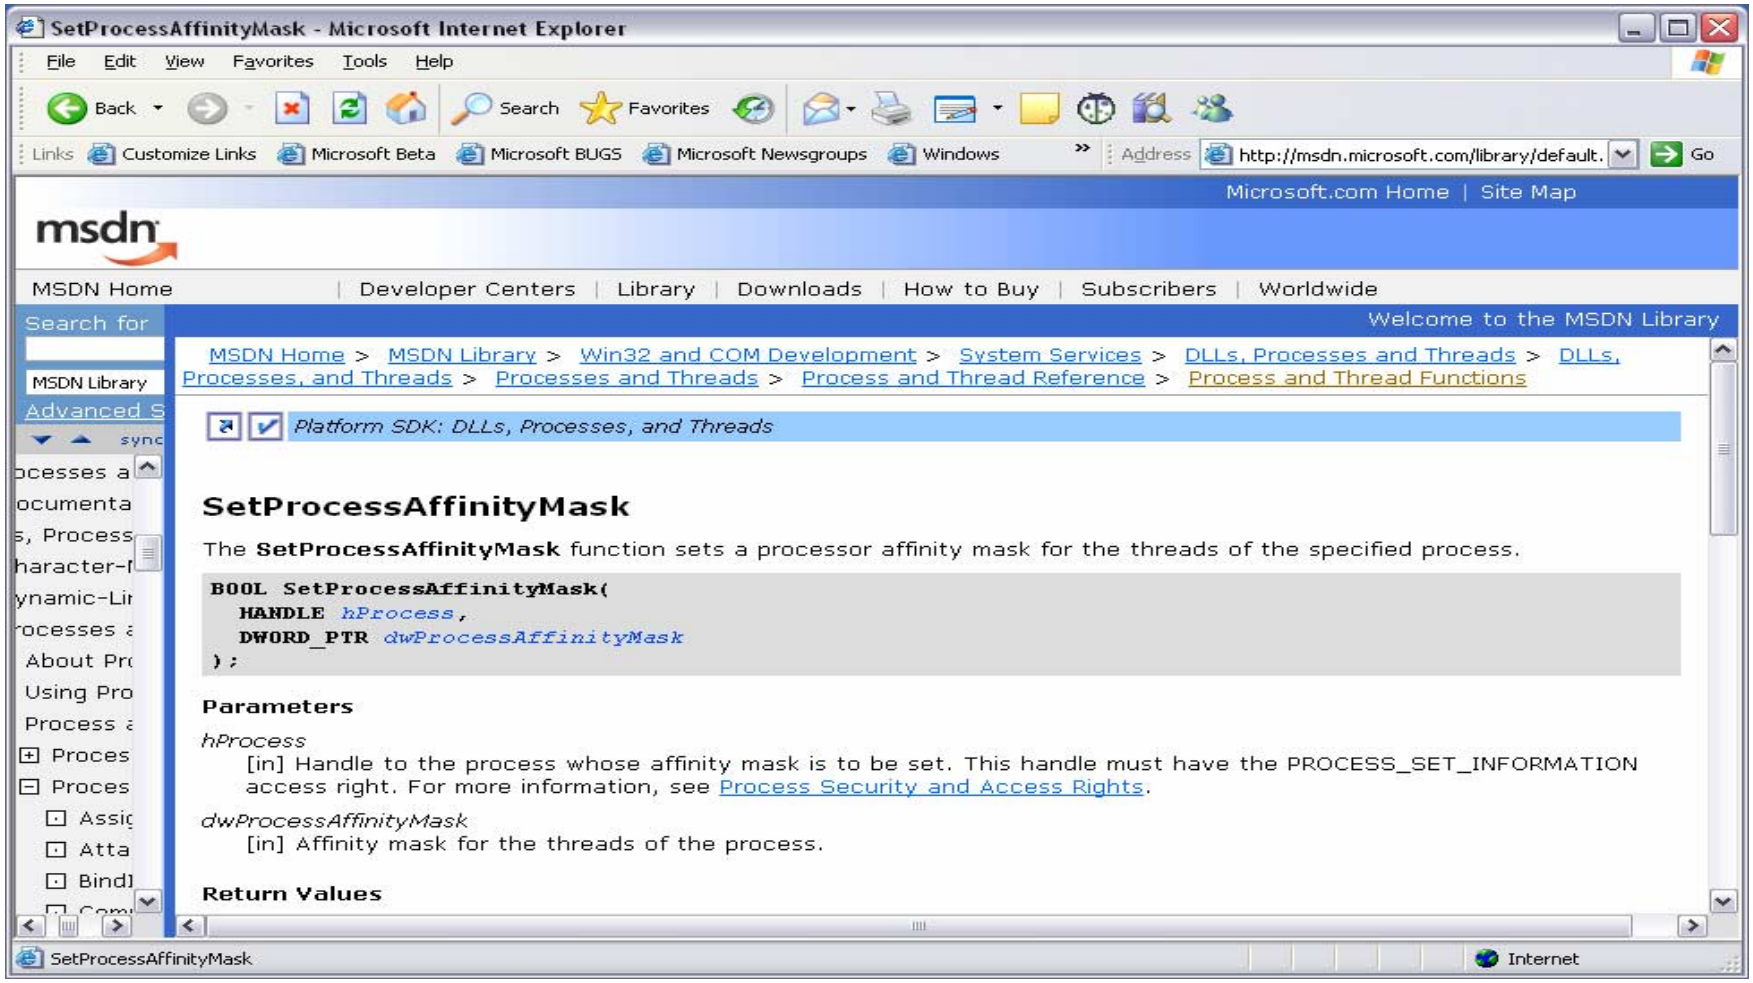

http://msdn.microsoft.com/library/default.asp?url=/library/en-us/dllproc/base/setprocessaffinitymask.asp

Windows® and AMD64 Technology

Unifying theme: Compatibility

- Processor: Native hardware support for 32-bit and 64-bit x86 code
- OS: 64-bit Windows® runs 32-bit and 64-bit applications side by side, seamlessly
- Code: A single C/C++ source code tree compiles to both 32-bit and 64-bit binaries

AMD64 Programmer's Model

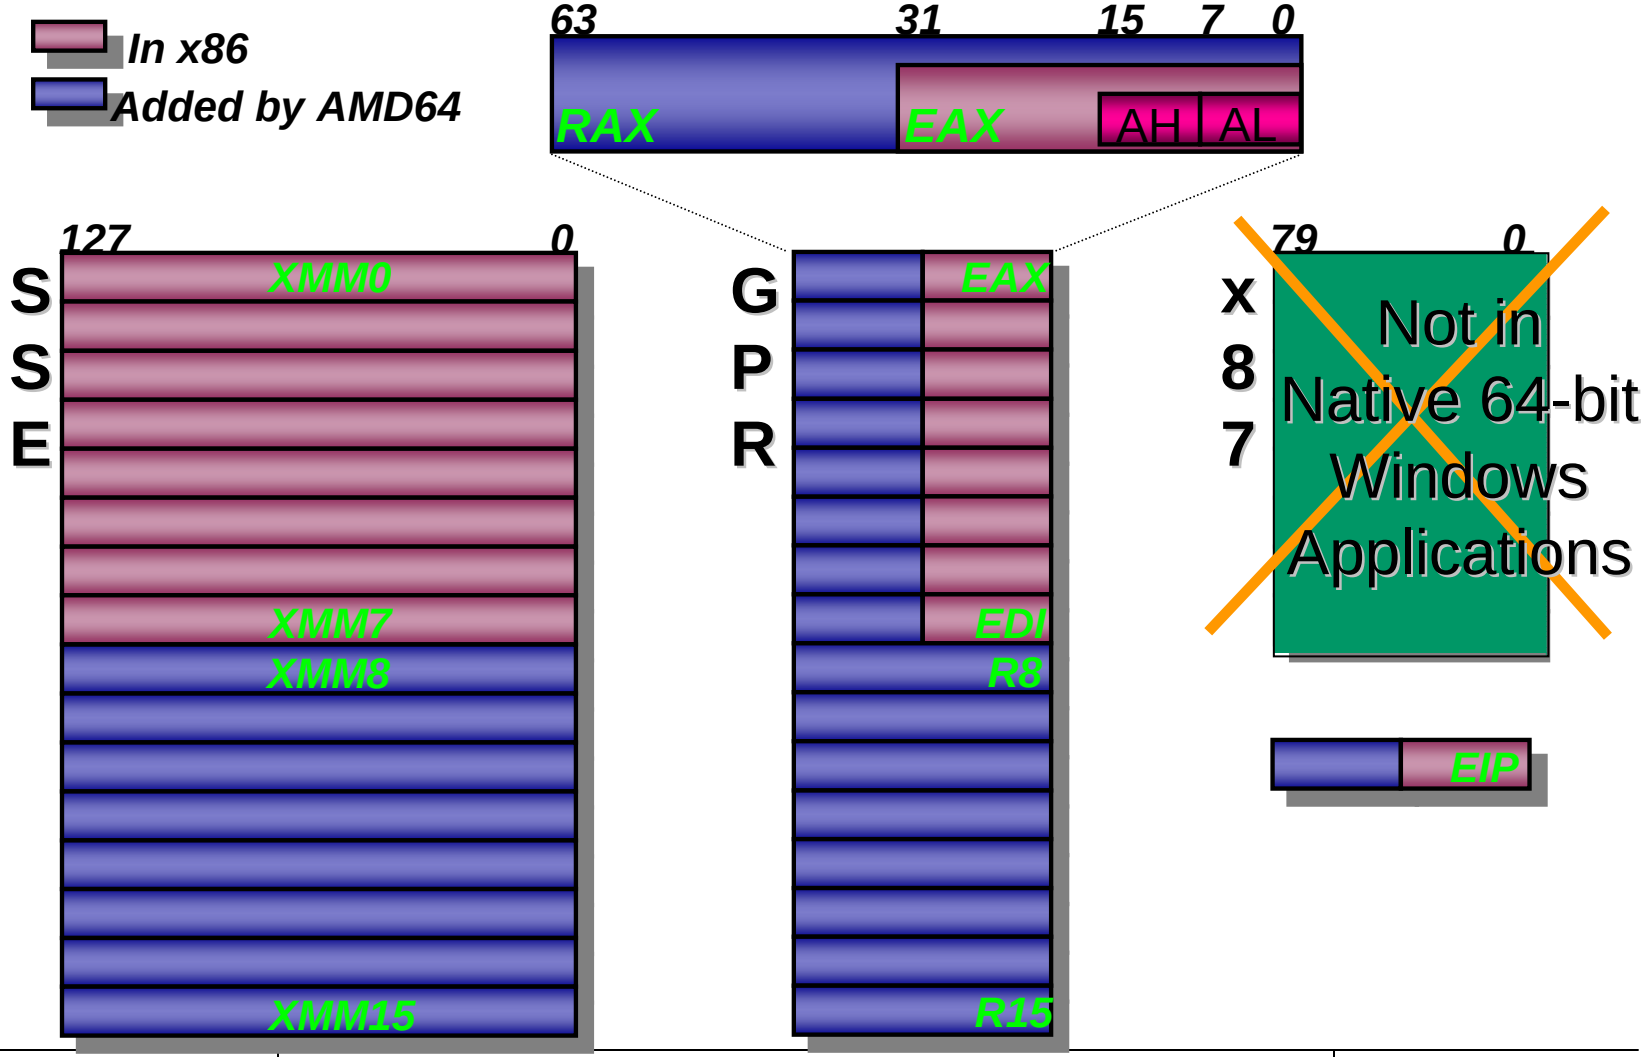

AMD64 Instruction Set Architecture

- Support for all x86 instruction extensions
  - SSE, SSE2, MMX™, x87, 3DNow!™
- Full performance with all kinds of code
  - Native 32-bit x86 mode
  - Enhanced capability in 64-bit mode
    - 64-bit general purpose registers
    - 64-bit addressing
    - Twice as many general purpose registers
    - Twice as many SSE registers
- Same familiar x86 instructions

WoW64 operation and compatibility

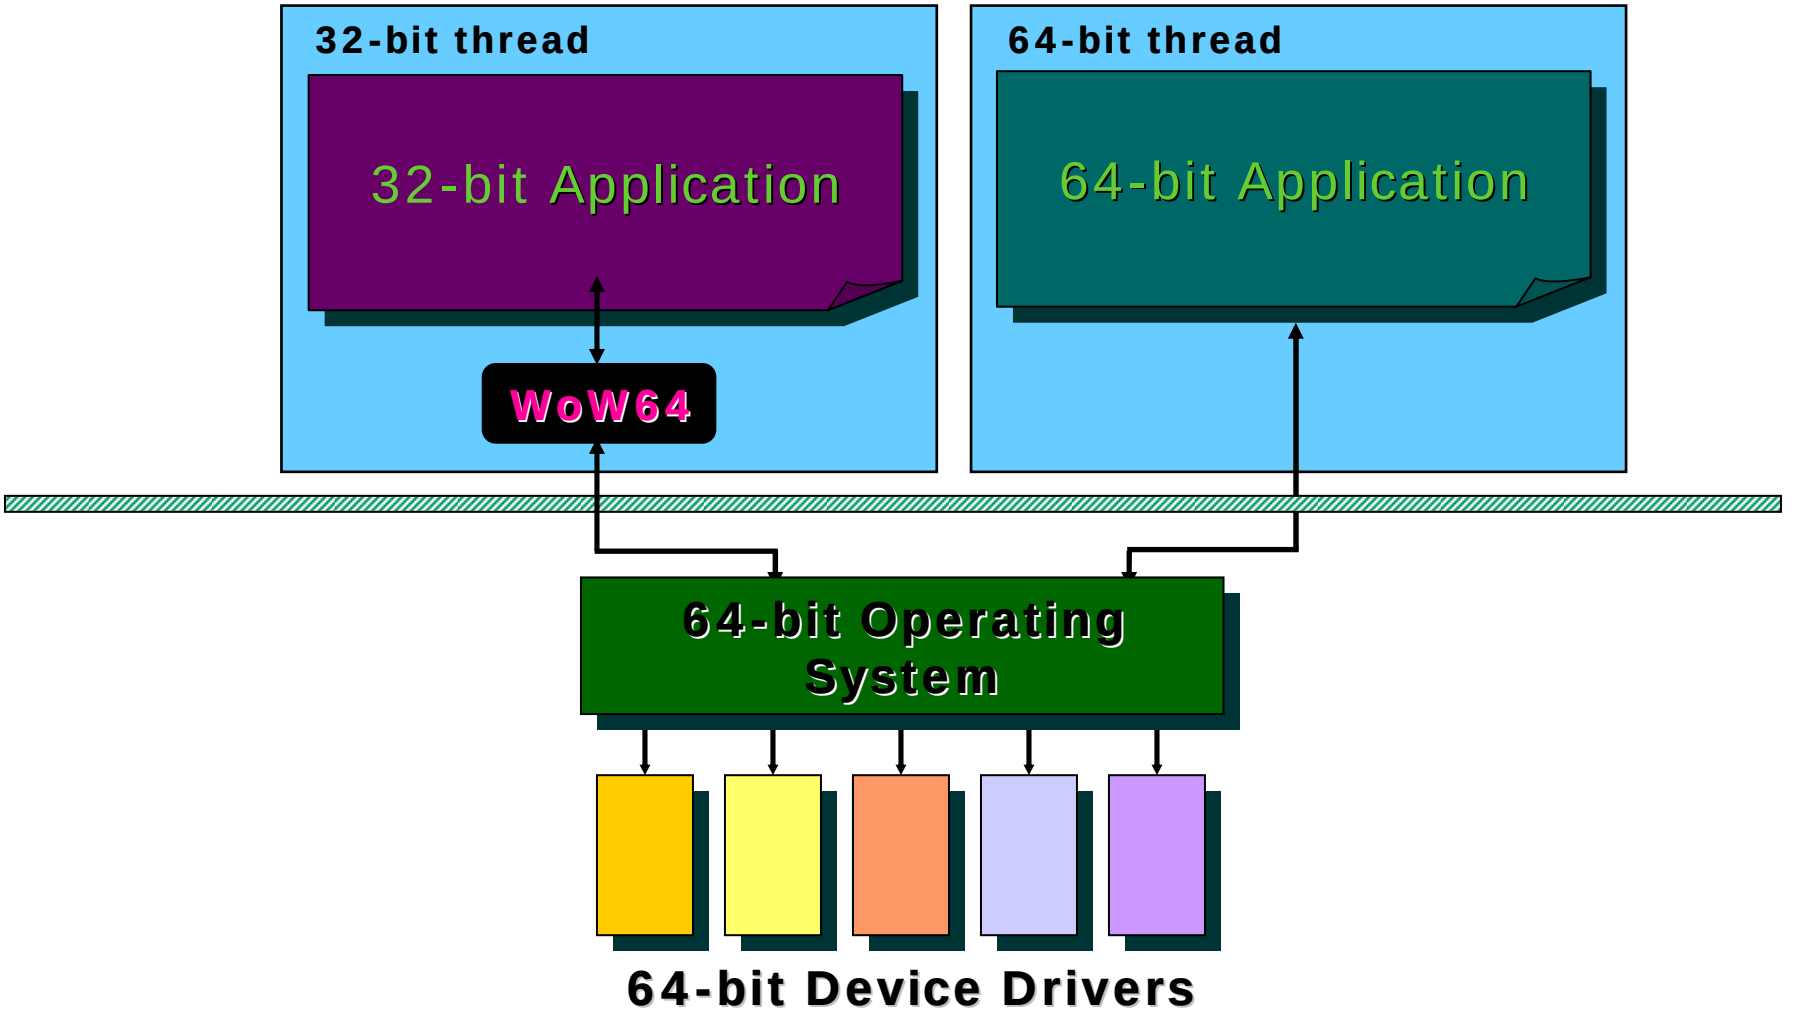

Why the 64-bit OS Benefits 32-bit Applications

32-bit x86 system

- OS and applications share virtual and physical memory
- Results in a lot of paging of info in and out of memory
- Limits the size of files and datasets

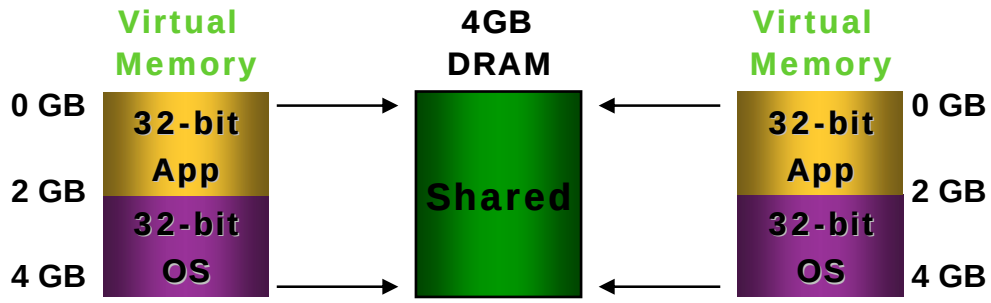

64-bit AMD64 system

- OS uses Virtual Memory space outside range of 32-bits
- Application has exclusive use of virtual memory space
- OS can allocate each application dedicated portions of physical memory
- Reduces the amount of paging
- Supports larger file and dataset sizes

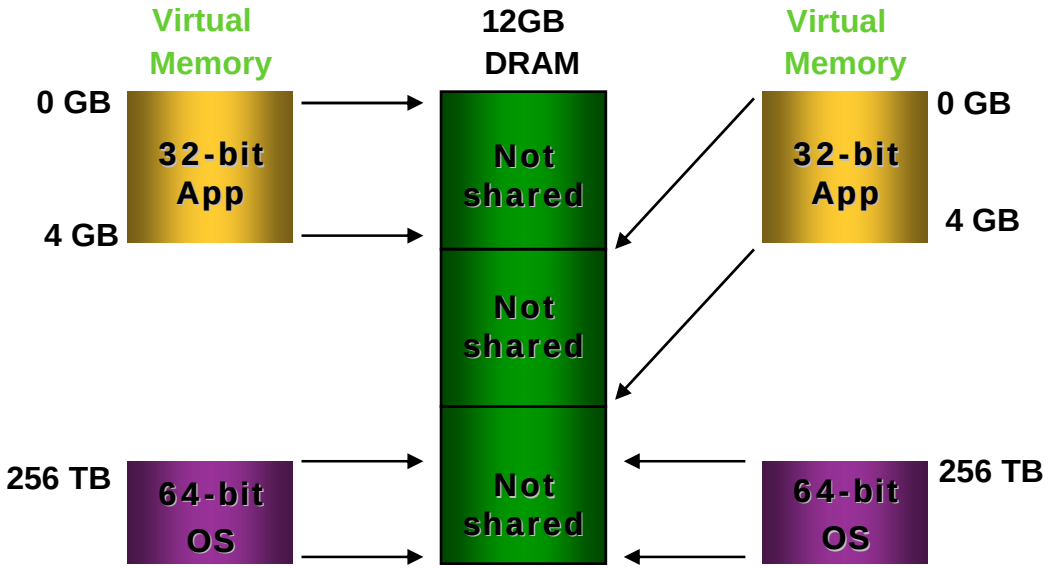

/LARGEADDRESSAWARE

- Use this linker switch to tell the linker that your 32-bit application can handle addresses larger than 2 gigabytes.
- This will work for most code, if your code takes advantage of the 2GB limit, then it will have to fixed.
- In Visual Studio, this flag is in the Property Page, under Linker->System->Enable Large Addresses.

AMD64: What to port

What apps will benefit from 64 bits?

AMD

- Large memory!
  - essentially unlimited virtual address space
  - physical memory only limited by platform capability
  - 8GB per CPU is expected to be common this year
- More registers and “big number” math
  - Codecs, simulation, 3D, games
  - Compression, encryption, finance

AMD64: What to port, and why

Some code really *must* be ported

- Drivers
  - device drivers must “match the OS”
  - 64-bit OS requires 64-bit drivers
  - there are presentations focused on drivers
- Code libraries and .dll’s
  - 64-bit applications will require 64-bit versions
- You can’t mix 32 and 64-bit application code
  - but OS IPC mechanisms work 32 ←→ 64

AMD64: How to port

Assorted portability tips

- -1 != 0xFFFFFFFF (just use “-1”)
- Many Windows® apps use DWORD
  - a DWORD is always 32 bits, make sure that is really what you want
- Use %p in a printf, to print a pointer
- Data alignment can affect performance
  - there are alignment requirements in the ABI
  - structure padding can affect portability
- Be watchful for unsigned int’s as array indices

Windows data types in 64-bits

pointers	Are always 64-bits in 64-bit code
int	Integers are still 32-bits
unsigned int	Of course is also 32-bit
char	8 bits
long	32-bits
long long	64-bits
float	32-bits
double	64-bits
long double	64-bits (but is considered a different type than ‘double’)

Some new data types in 64-bit Windows

Fixed Precision

INT32	32-bit signed integer
INT64	64-bit signed integer
DWORD32	32-bit unsigned integer
DWORD64	64-bit signed integer
LONG32	32-bit signed integer
LONG64	64-bit signed integer

AMD

Pointer Precision

INT_PTR	signed integer of pointer precision
UINT_PTR	unsigned integer of pointer precision
LONG_PTR	signed long of pointer precision
ULONG_PTR	Unsigned long of pointer precision

New Specific Pointer Precision Types

Specific Pointer Precision

POINTER_32	A 32-bit pointer
POINTER_64	A 64-bit pointer

Be careful with these. For ***64-bit*** Windows, POINTER_32 will just be a truncated 64-bit pointer.

For ***32-bit*** Windows, POINTER_64 will be a sign-extended 32-bit pointer.

Microsoft C compiler and libraries

- Microsoft 64-bit compiler for AMD64
  - SSE and SSE2 code always used for floating point
  - use optimization switches: /O2, /fp:fast (careful with this one)
  - use /GL for Whole Program Optimization (WPO)
  - use Profile Guided Optimization (PGO)
    - Available in Whidbey
- The libc string and memory functions are optimized and should be used
- AMD provides optimized math libraries:
  - ACML = AMD Core Math Libraries (BLAS, FFT, LAPACK, and more)

What's Wrong With This Picture?

AMD

```
#include <stdlib.h>
#include <stdio.h>

int main() {
   int            x  =  -30;
   unsigned int   y  =   20;
   int            *p1, *p2;
   p1 = new int[50];
   for (int i=0; i<50; ++i)
      p1[i] = i;
   p2 = &p1[25];
   p2 = p2 + (x + y);

   printf("%d\n", *p2);
   delete[]p1;
Return 0; }
```

MSFT Community Resources

- Community Resources

  http://www.microsoft.com/communities/default.mspx

- MSDN Newsgroups

  Converse online with Microsoft Newsgroups, including Worldwide

  http://msdn.microsoft.com/newsgroups/default.aspx

- User Groups

  Meet and learn with your peers

  http://www.microsoft.com/communities/usergroups/default.mspx

More Resources - AMD

- AMD Developer Center
  - come to California, use our machines, talk with our tech people
  - remote access via VPN
  - https://devcenter.amd.com/
- AMD tools and documentation
  - CodeAnalyst profiler, Optimization Guide, Programmer's Manuals, other tech docs.
  - AMD64 Developer Resource Kit
- Go to http://www.amd.com "Develop with AMD"

Introducing AMD64 Dual Core Processor

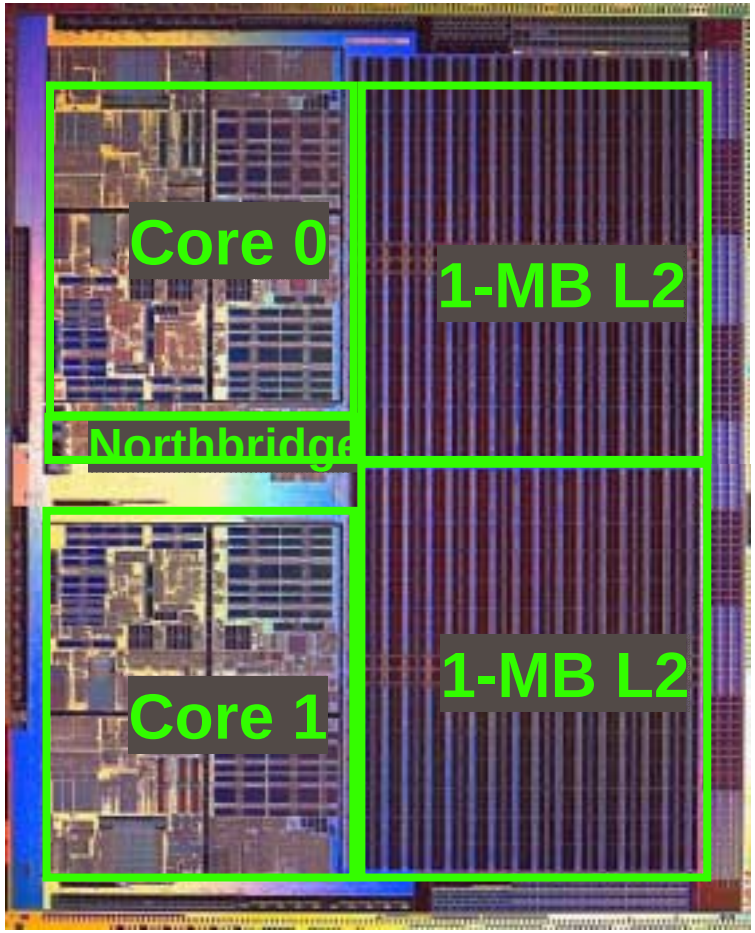

- **Two AMD Opteron™ CPU cores on one single die, each with 1MB L2 cache**
- **90nm, ~205 million transistors***
  - **Approximately same die size as 130nm single-core AMD Opteron processor***
- **95 watt power envelope fits into 90nm power infrastructure**
- **940 Socket compatible**
  - **AMD expects to be first to introduce dual-core for the one- to eight-processor server and workstation market in mid-2005.**
- **Dual-core processors for client market are expected to follow beginning in 2H'05.**

Based on current revisions of the design

Dual-Core AMD Opteron™ Architecture

- Designed for Dual Core From the Start
- Shared Northbridge 3 HyperTransport™ technology links
- Two complete CPUs
- Non-disruptive migration

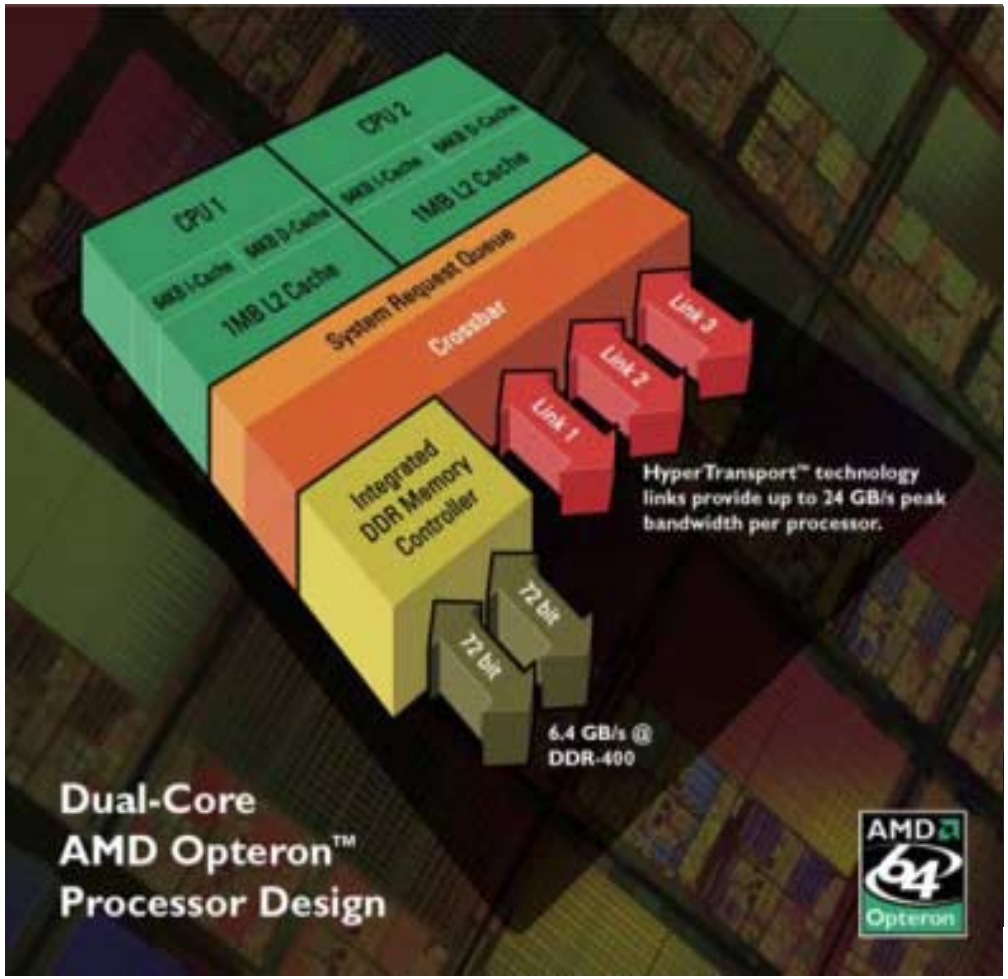

Operating System Support for Dual-Core

- Existing OS kernels support AMD Dual-Core using CMP Legacy Mode.
  - First AMD dual-core silicon using this model booted Windows® and Linux within hours.
  - Recent OS distributions that support Hyperthreading are expected to work well.

- New extended CPUID function (eax=8000_0008) returns on any core the number of physical cores per processor
  - Correct way for future OS and Application software

General Observations Regarding Dual-Core and Threading

- Software should use OS APIs to test for number of processors
  - Software may also use CPUID functions to specifically test for Dual-Core capability.

- Threading Performance in Multi-Processor systems also affected by OS Memory Affinity Algorithms
  - May require threads to allocate their own memory at thread initialization rather than relying on parent or the heap.
    - Linux for AMD64 assigns memory affinity at first use rather than allocation; other Operating Systems may be different.
  - Not an issue for a single-processor Dual-Core system.

AMD, the AMD Arrow Logo, AMD Athlon, AMD Opteron, and combinations thereof, 3DNow!, AMD-8111, AMD-8131, and AMD-8151 are trademarks of Advanced Micro Devices Inc. HyperTransport is a licensed trademark of the HyperTransport Technology Consortium. Windows is a registered trademark of Microsoft Corporation. MMX is a trademark of Intel Corporation. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.