

Microsoft®

Tech·Ed
2005

Learn · Solve · Grow

AMD64 with Direct Connect Architecture: Providing a Scalable Platform for Microsoft SQL Server

Name: Margaret Lewis

Title: Commercial Software Strategist

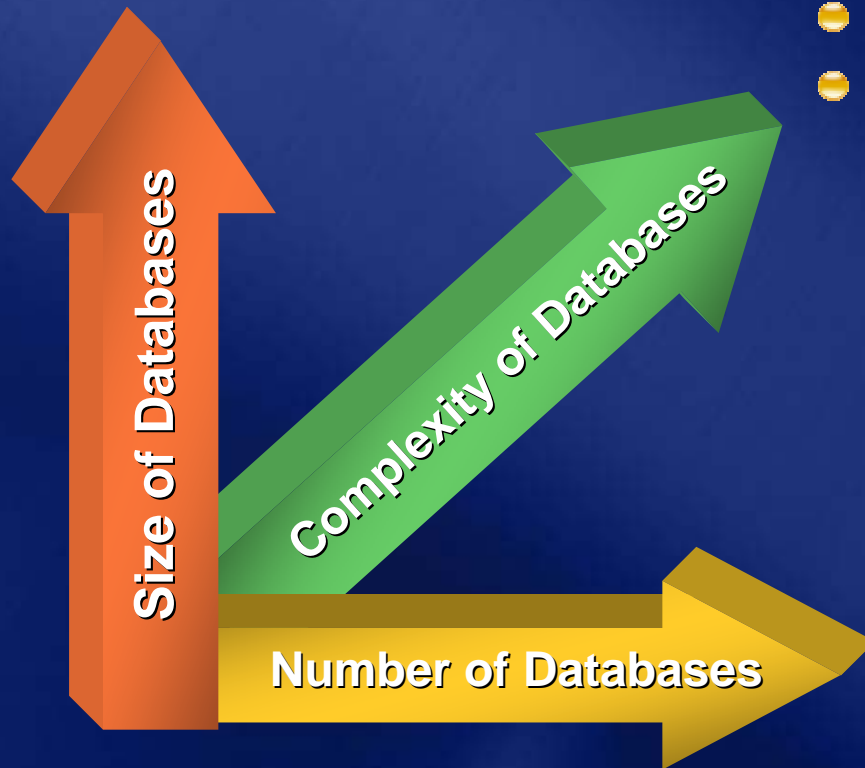
Company: AMD

Database Growth

A Three Dimensional Problem

- New Business Data
- Historical Data

- Multiple Database Formats
- Bridging Multiple Databases



- Business Intelligence
- CRM
- Ecommerce
- Enterprise Portals

Database Server Requirements

- Supply users with fast access to information
- Deliver performance that scales as demands increase
- Keep pace with expanding data stores and file sizes
- Provide a secure, reliable platform that fits within infrastructure and cost constraints



Achieving Scalability

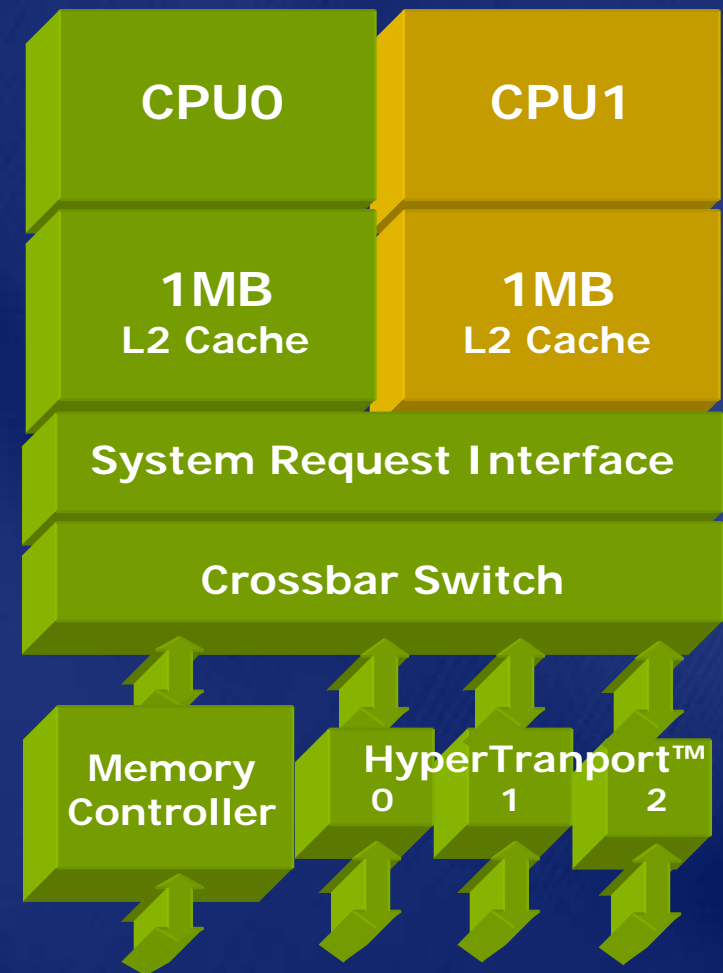
- Ability of hardware and software to increase or decrease capacity to support larger or smaller volumes of data and more or less users
- Factors influencing scalability
 - 64-bit capabilities
 - Multi-core processors
 - Fast access to memory
 - High-speed I/O



Achieving Scalability

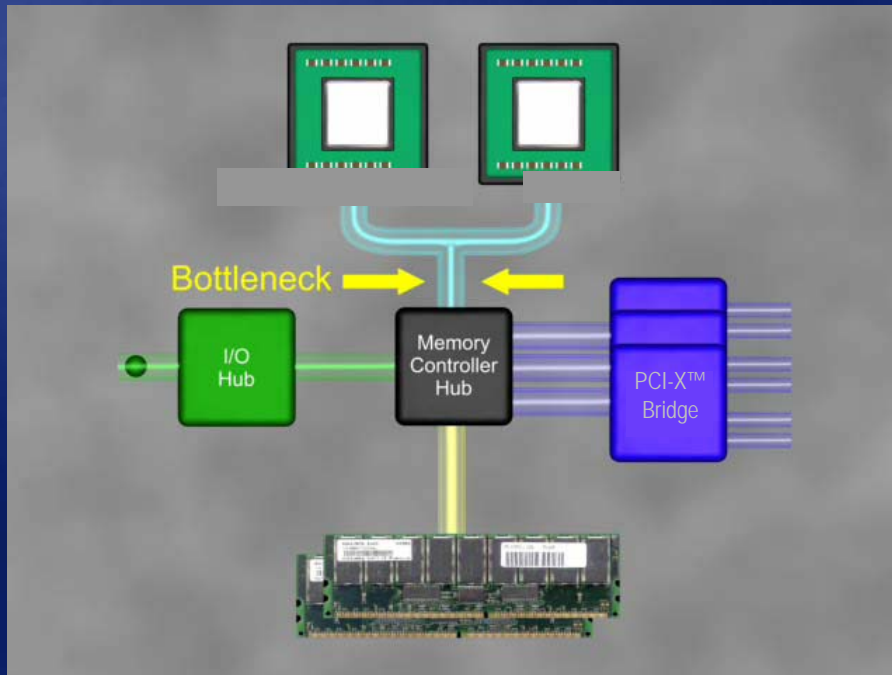
AMD Opteron™ Processor with Direct Connect Architecture

- **AMD64 Core**
 - Enables 32- and 64-bit computing
 - Eliminates 4GB memory barrier
 - Dynamically throttled core
- **Integrated Memory Controller**
 - Reduces memory latency
- **HyperTransport™ Technology**
 - Up to 24.0 GB's peak bandwidth per processor
 - Directly connects CPUs




Achieving Scalability

Front-side Bus



- Traditional platform
- Processors on a shared bus
- Same bus for memory & I/O
- 64-bit afterthought

Direct Connect Architecture



- Architecture for the Future
- Direct Connect for processors
- Direct Connect for memory & I/O
- Designed for 64-bit and 32-bit

Achieving Scalability



- **Support for greater memory space**
- **Performance improvements due to**
 - **Improved memory management**
 - **Expanded registers**
 - **I/O subsystems**
- **Better protection from malicious code**
- **Compatibility with existing 32-bit applications**

Achieving Scalability

Microsoft SQL Server 2005

Advantages of transitioning from 32-bit to 64-bit database


- **Significant portions of a database can be loaded into memory**
- **Large virtual address provides more space to efficiently process data**
- **Information for a large data set can be stored in a single file**

64-bit Capabilities

64-bit Capabilities

AMD Opteron™ Processor

- **Processes more data per clock cycle**
 - 64-bit wide registers
 - Performs operations on 64 bits of data at one time
- **Addresses more memory**
 - 256TB (48 bits) of virtual memory
 - 1TB (40 bits) of physical memory



64-bit Capabilities

Windows® Server 2003 x64

- **Eliminates 4GB virtual memory limitation of 32-bit OS**
 - **By default 2 GB is addressable by a process and 2 GB addressable by 32-bit Windows®**
 - **Can be reconfigured to give 3 GB to a process and 1 GB for Windows - but this squeezes the OS!**

64-bit Capabilities

Windows® Server 2003 x64

General Memory Limits	32-Bit Windows®	64-Bit Windows
Total virtual address space	4 GB	16 TB
Virtual address space per 32-bit app	2 GB (3GB if system is booted with /3GB switch)	4 GB if compiled with /LARGEADDRESSAWARE
Virtual address space per 64-bit app	Not applicable	8 TB
Paged pool	47 MB	128 GB
Non-paged pool	256 MB	128 GB
System Page Table Entry (PTE)	660 MB to 900 MB	128 GB

Physical Memory & Processor Limits	32-Bit Versions	64-Bit Versions
Windows XP Professional	4 GB / 1 to 2 Processors	128 GB / 1 to 2 Processors
Windows Server 2003, Standard Edition	4 GB / 1 to 4 Processors	32 GB / 1 to 4 Processors
Windows Server 2003, Enterprise Edition	64 GB / 1 to 8 Processors	1 TB / 1 to 8 Processors
Windows Server 2003, Datacenter Edition	64 GB / 8 to 32 Processors	1 TB / 8 to 64 Processors

64-bit Capabilities

Microsoft SQL Server 2005

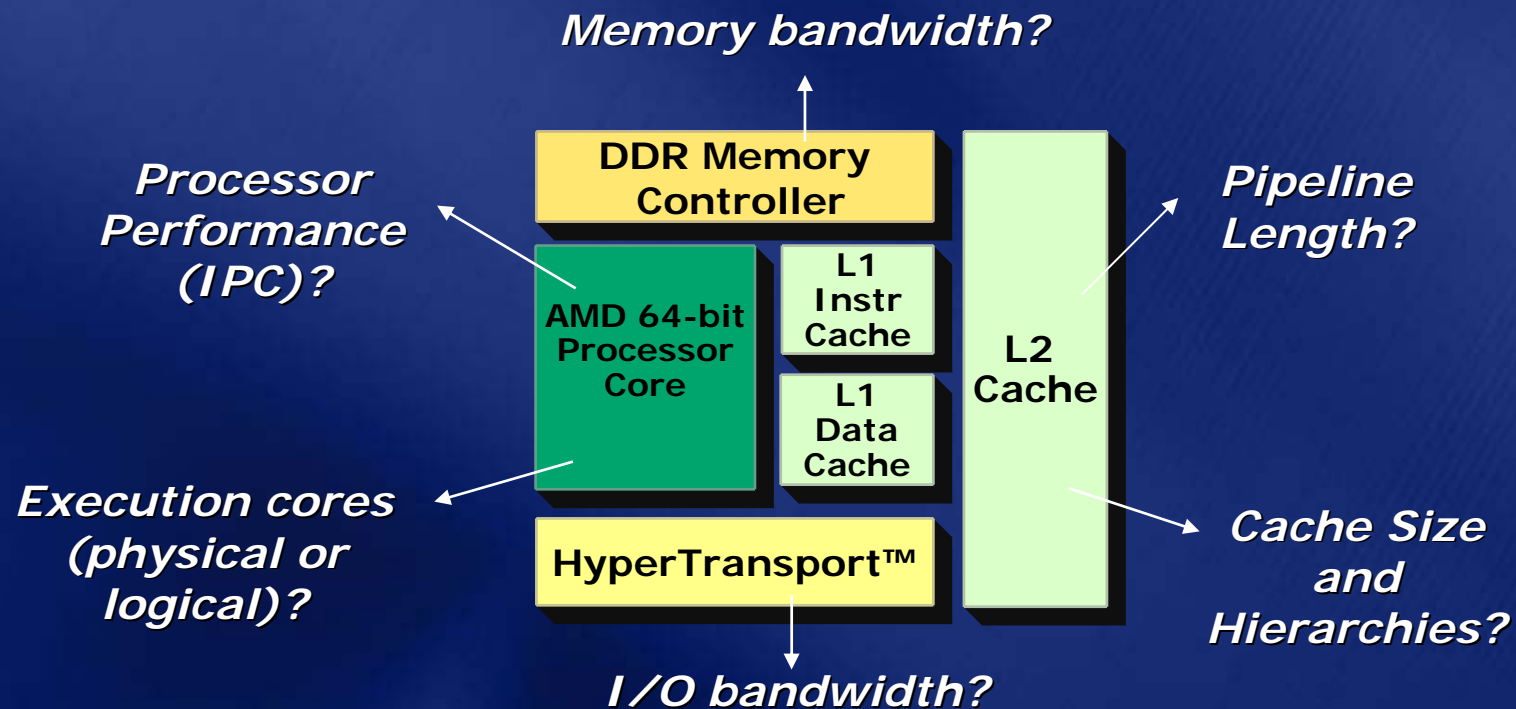
Examples of database operations that benefit from 64-bit:

- Queries that sort large volumes of data
- Complex, multi-step queries
- Multiple database objects are defined and open simultaneously
- Handling many concurrent users, cursors, transactions, locks, etc.
- Handling many stored procedures

Multi-core Processors

Why Multi-Core Processors?


Size, power consumption, and performance needs of today's computers require new innovations in processor design



Multi-Core Processors

AMD Opteron™ Dual-Core Technology

- Designed for dual core
- 940-pin compatible
- No changes in power
- Non-disruptive migration

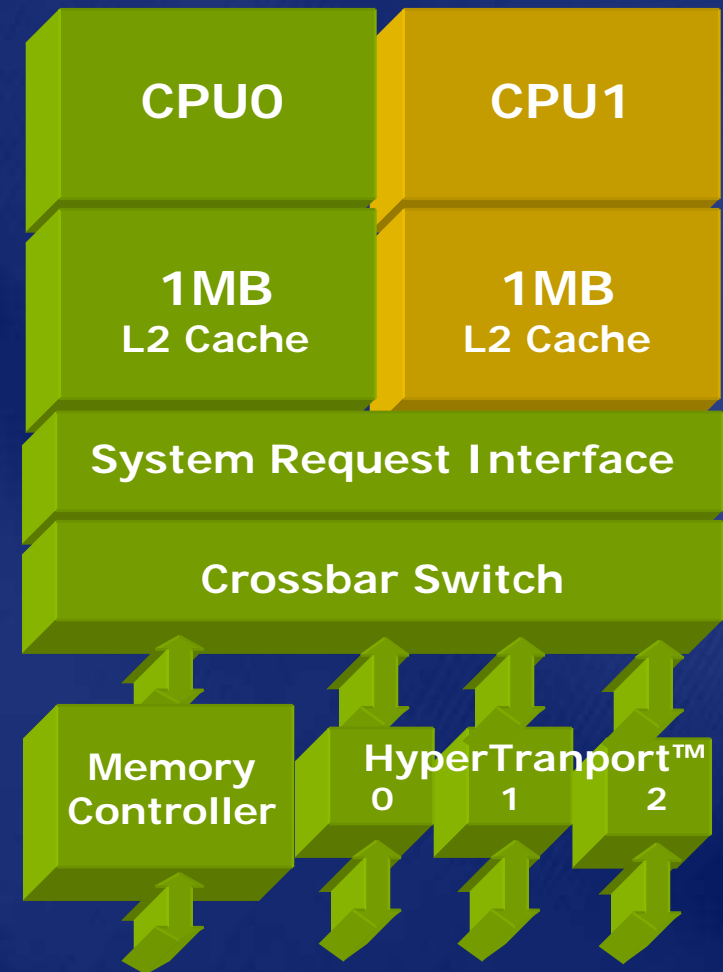


Photograph of 90nm dual-core
AMD Opteron processor prototype

Multi-Core Processors


AMD Opteron™ Dual-Core Technology

- One die with 2 CPU cores
- Each core has separate L1/L2 cache hierarchies
- Cores share Integrated Memory Controller and HyperTransport™ links
- Supports AMD PowerNow!™ Technology with Optimized Power Management



Multi-Core Processors

Dual-Core Performance




SPEC and the benchmark names SPECint and SPECfp are registered trademarks of the Standard Performance Evaluation Corporation. SPEC score for AMD Opteron processor Model 265, 275, 865 and 875-based systems under submission to SPEC as of April 18, 2005. Other competitive results stated above reflect results published on www.spec.org as of April 6, 2005. For the latest SPECint and SPECfp results visit <http://www.spec.org/cpu2000>.

Dual-core technology allows AMD to offer a competitive performance roadmap while meeting the power requirements of our customers

Fast Access to Memory


Fast Access to Memory

Memory Controller Technology



Integrated Memory Controller

- Runs at CPU frequency
- CPU has dedicated bandwidth for memory
- Bandwidth increases with addition of CPU's
- No additional power required

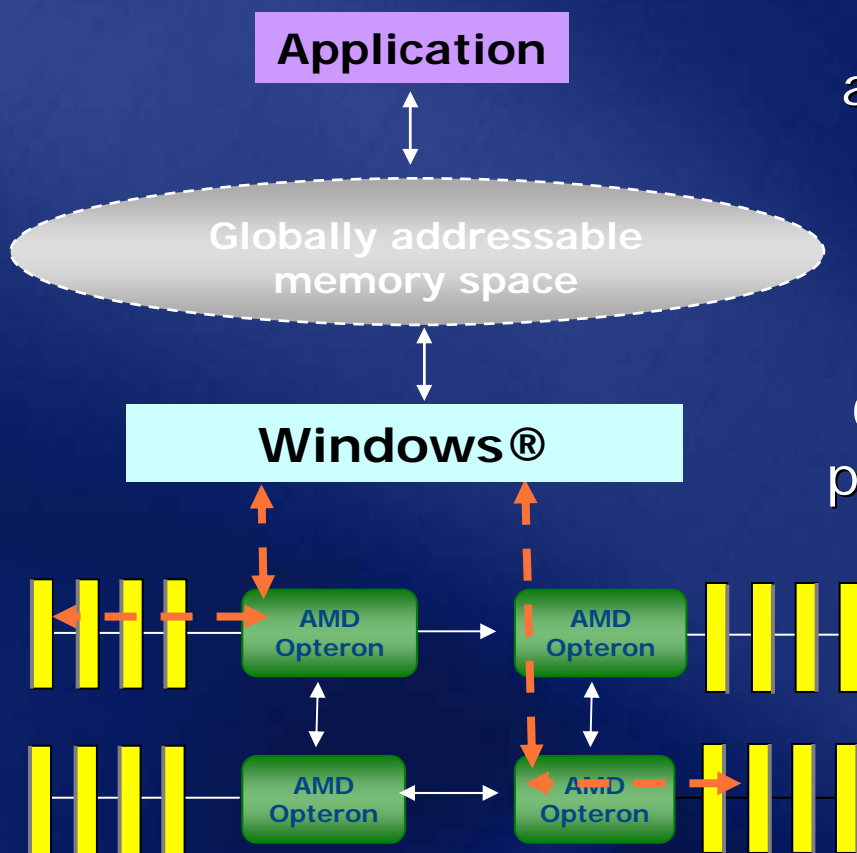


External Memory Controller

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

Fast Memory Access

Non-Uniform Memory Access (NUMA)




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

OS assigns threads from same process to the same NUMA node

Processor has local memory and uses HyperTransport™ for high-speed access to non-local memory

Fast Access to Memory

Memory Bandwidth



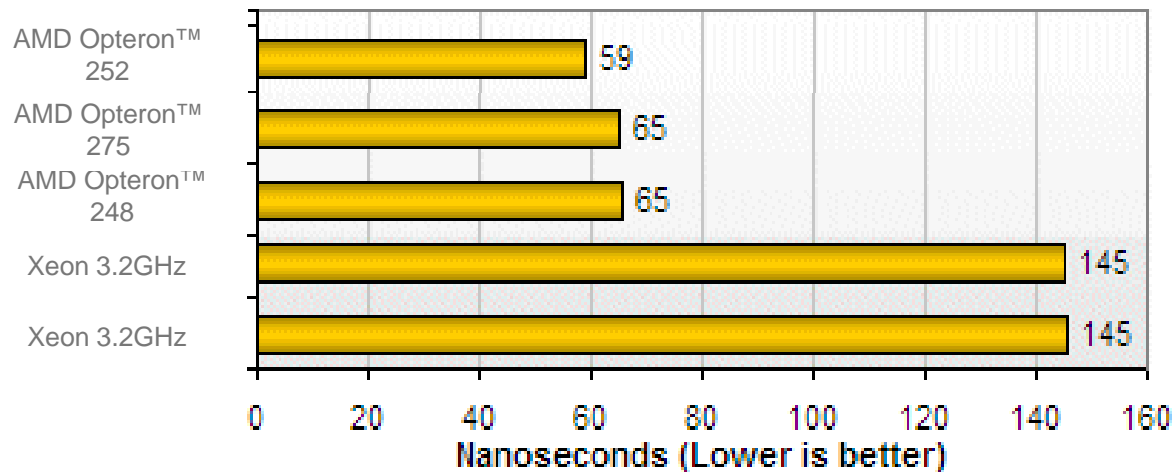
Memory
bandwidth
increases
with
frequency

SiSoft Sandra is a synthetic test for stressing individual components. The tests can stress a system's CPU, Memory, or Multimedia capabilities. Higher Sandra scores mean better overall component performance.

Fast Access to Memory

Memory Latency

Scencemark - Memory Access Latency
8192KB block, 512-byte stride



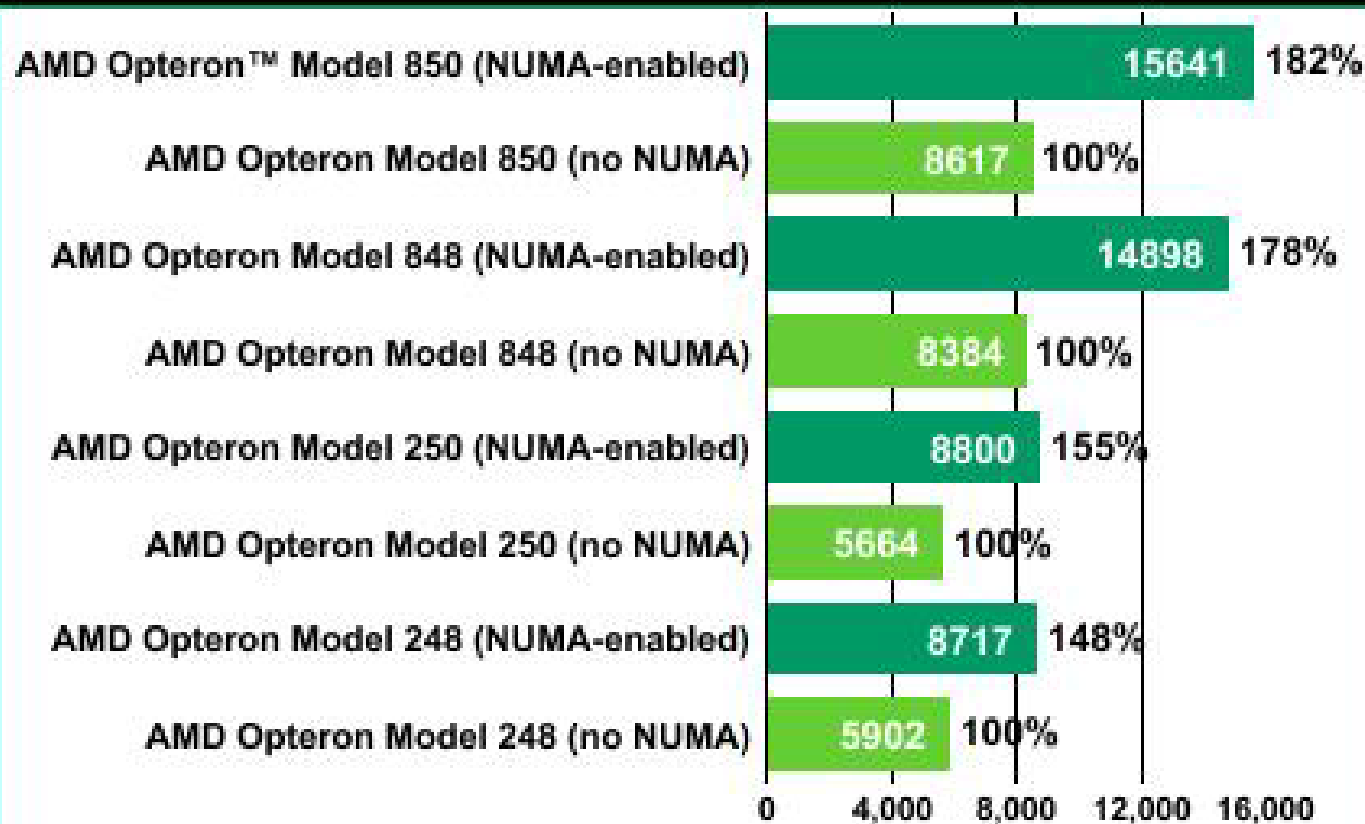
**Memory
latency
decreases
with
frequency**

Scencemark is a synthetic benchmark which stresses CPU performance. The benchmarks are a collection of science and physics mathematical calculations. Lower times mean better performance.

Fast Access to Memory

NUMA Performance

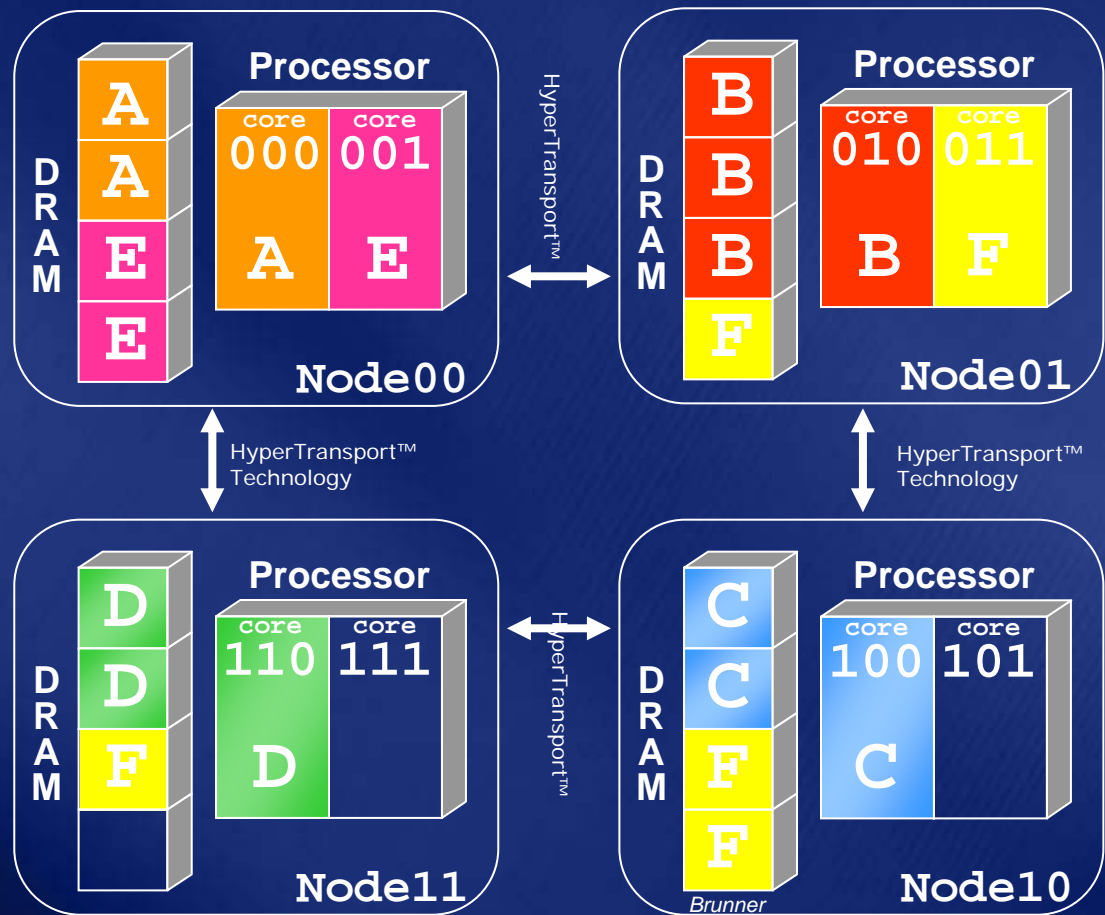
NUMA vs. Non-NUMA Performance (using STREAM)



Fast Memory Access

NUMA Scheduling Example

- OS distributes threads across processors
- OS attempts to map a thread's memory requests to local memory
- If no memory is left on the current node, OS will use memory of other nodes



High-speed I/O

High-Speed I/O

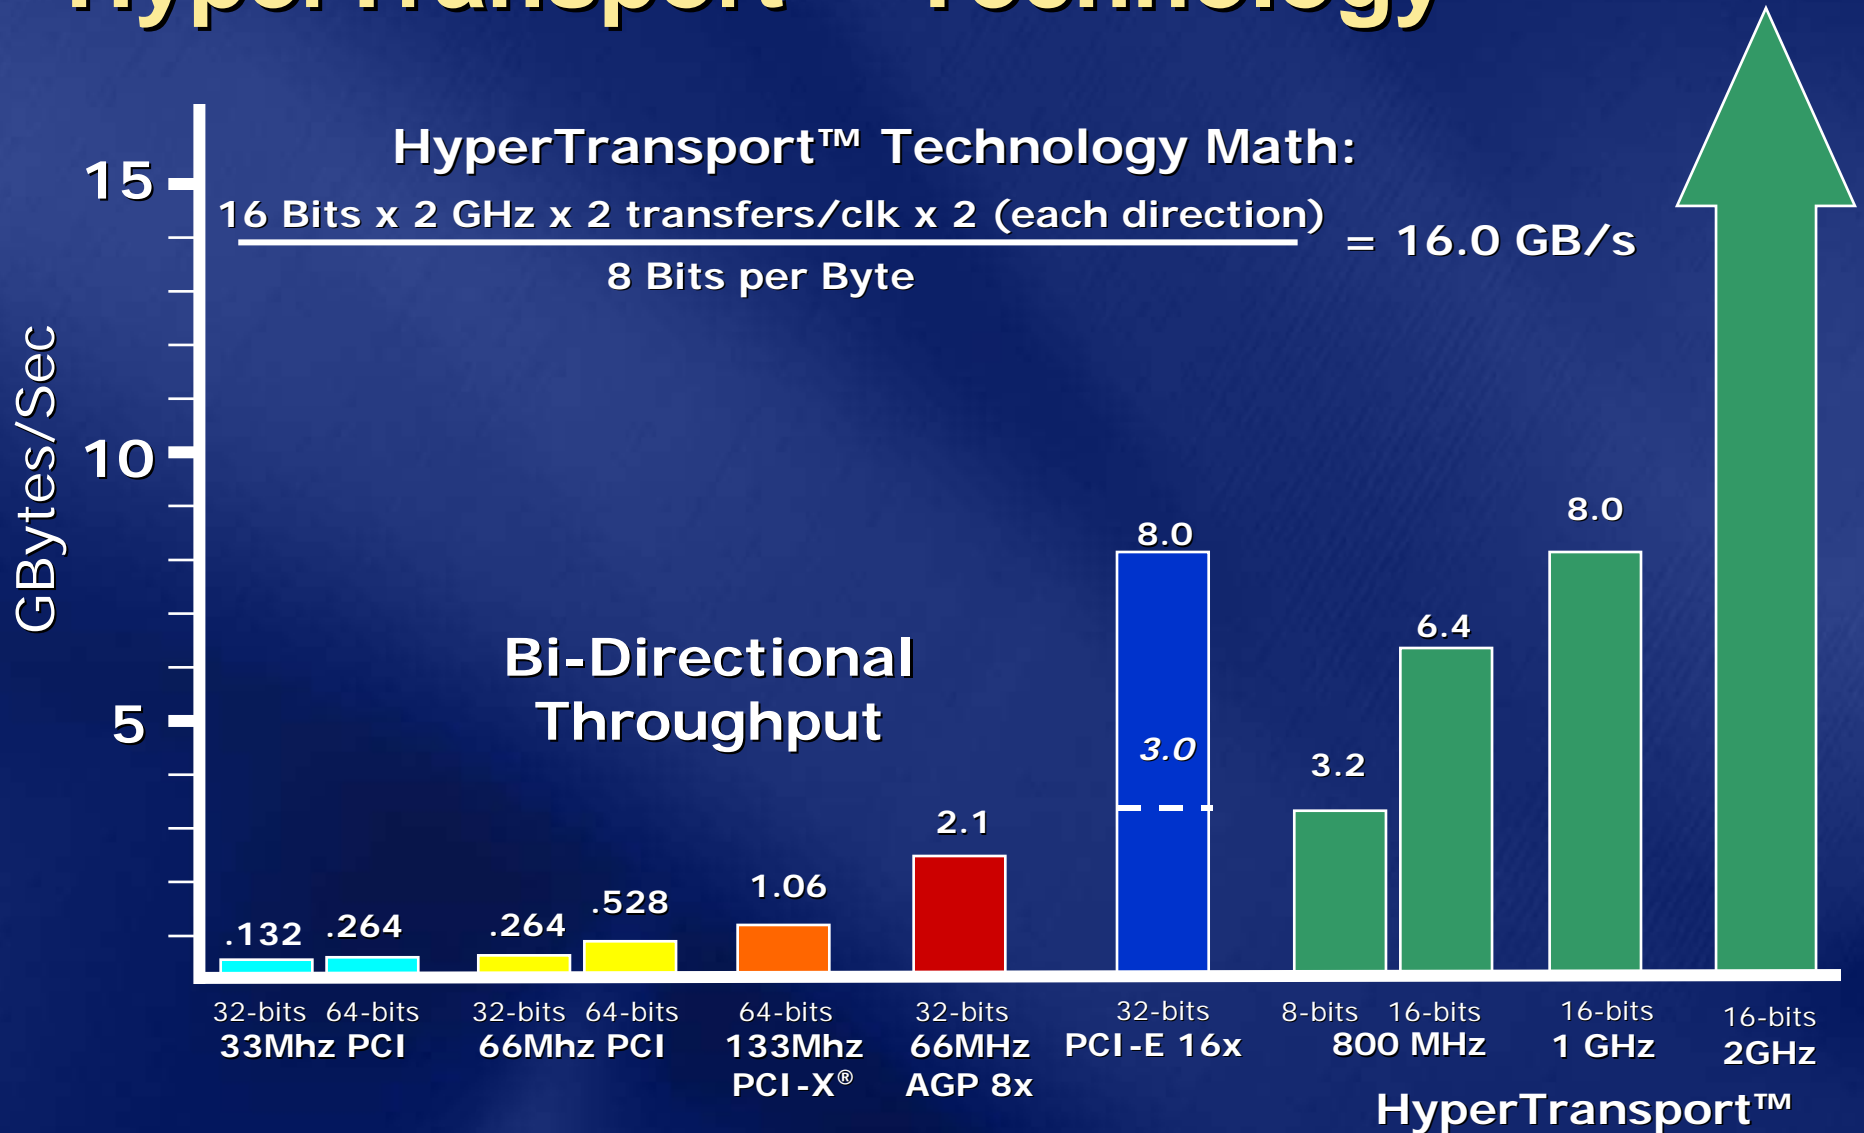
HyperTransport™ Technology

- High performance chip-to-chip interface
- Full PCI software compatibility
- Supports up to 31 devices
- Allows building block approach to system design
- Reduces the number of proprietary buses
- Open standard
- Cost efficient



High-Speed I/O

HyperTransport™ Technology



High-Speed I/O


HyperTransport™ Technology

Typical Communication Latencies

Ethernet	>10 μ S
Infiniband via PCI-X	4 to 10 μ S
High-Speed Fabric via PCI-X	2 to 4 μ S
HyperTransport™ Connected	< 2 μ S

High-Speed I/O

HyperTransport™ Technology




- Enables “glueless” multi-CPU design
- Each CPU adds HyperTransport™ bandwidth
- Three (coherent capable) links per AMD Opteron™ 800 Series
- Scales up to 8 processors/16 cores

Putting the Pieces Together

Putting the Pieces Together

TPC-C Benchmark




Running Windows® Server 2003 Enterprise Edition
and Microsoft SQL Server 2000

Putting the Pieces Together

SAP SD 2-Tier Benchmark

SAP Standard Application Sales and Distribution (SD) 2-Tier Performance
Microsoft® Windows® Server 2003 Enterprise Edition



- Scales from single core to dual core
- Lower power consumption
- Supports 32-bit and 64-bit applications

Microsoft and Windows are registered trademarks of the Microsoft Corporation in the United States and other jurisdictions. For complete SAP benchmark information visit <http://www.sap.com/benchmark>.

Running Windows® Server 2003 x64 Enterprise Edition
Microsoft SQL Server 2000

http://www.amd.com/us-en/Processors/ProductInformation/0,,30_118_8796_8800~97419,00.html

Putting the Pieces Together

Microsoft SQL Server 2005 WinHEC Demo



- Compares performance of 32-bit and 64-bit database environments
- Loads 5 million records into look-up table and runs client queries
- 64-bit Microsoft SQL Server 2005 running a dual-core HP ProLiant DL585 handled 5 times the queries of 32-bit configuration

Putting the Pieces Together

Microsoft SQL Server 2005 Case Study

AMD Opteron-based server + Microsoft SQL Server 2005
Performance to handle large amounts of data while providing users quick response to their queries

Problem: Microsoft's Volume Licensing Business Tools and Analytics group needed more effective way to gather pricing data and generate enterprise agreements

Solution: Microsoft Office smart client application, "the Easy Pricing and Estimation Tool" (EZPET 2.0)

- EZPET 2.0 allows users to select prices, find attributes that define the price of a product, and automatically prepare a quote
- Built on the Microsoft .NET Framework 2.0 using the Microsoft Visual Studio® .NET 2005 Beta development environment
- Pulls pricing data from databases with 3.9 million price points
- Data Transformation Service Wizard in SQL Server 2005 used to create smaller, manageable sets of data that can be delivered to the client tier

Achieving Scalability

Closing Thoughts

- 64-bit environment
 - Start the migration today!
- Multi-core technology
 - Performance without increasing power
- Innovative system design
 - Fast memory and I/O access that translates into better application performance



Resources

Resource 1:

http://www.amd.com/us-en/Processors/ProductInformation/0,,30_118_8796,00.html

Resource 2:

http://www.amd.com/us-en/Processors/ProductInformation/0,,30_118_8796_8800,00.html?redir=CPOS13

Resource 3:

<http://www.microsoft.com/windowsserver2003/techinfo/overview/x64benefits.msp>

Resource 4:

<http://www.amd.com/amd64ecosystem>

***Your Feedback
is Important!***



***Please Fill Out a Survey for
This Session on CommNet***

Trademark Attribution

AMD, the AMD Arrow Logo, AMD Opteron and combinations thereof are trademarks of Advanced Micro Devices, Inc. HyperTransport is a licensed trademark of the HyperTransport Technology Consortium. Microsoft and Windows are registered trademarks of Microsoft Corporation in the U.S. and/or other jurisdictions. Other names are used for informational purposes only and may be trademarks of their respective owners.

Microsoft[®]

Your potential. Our passion.[™]

© 2005 Microsoft Corporation. All rights reserved.

This presentation is for informational purposes only. Microsoft makes no warranties, express or implied, in this summary.