IBM® DB2® Universal Database™ Software on the AMD Opteron™ Processor

David Kaplowitz
Advanced Micro Devices, Inc.
One AMD Place
Sunnyvale, CA 94088

June, 2003

Executive Summary

Relational Database Management Systems (RDBMS) have become integral parts of most business-critical applications. Databases are growing larger and larger each year with many of today's enterprise databases exceeding one Terabyte in size. To efficiently process such large amounts of data, the RDBMS requires access to large amounts of memory. As a rule, the more memory present on the system the better the database will perform. However, the large sizes of today's databases have exceeded the 4 GB addressing limits of today's 32-bit x86 systems.

The solution to this problem is to provide a flat 64-bit programming model. AMD has created the AMD64 Instruction Set Architecture (ISA) to support 64-bit addressing. The AMD64 ISA extends the industry-standard x86 architecture rather than creating an entirely new instruction set. IBM has recently ported the DB2 Universal Database product (DB2) to the AMD64 instruction set on the AMD Opteron™ processor as a native 64-bit application.

This white paper presents an overview of the importance of 64-bit addressing to database applications, the AMD Opteron processor, the DB2 software architecture, and a brief summary of the experience the DB2 software team had porting their application to the AMD64 instruction set running on the AMD Opteron processor.

Databases and 64-bit Addressing

Databases are a key component in today's business-critical enterprise applications such as ERP, CRM, SCM, and Business Intelligence. With enterprises wanting to automate and integrate more and more aspects of their business, databases have consistently been growing larger and larger in size. Today's largest On-Line Transaction Processing (OLTP) databases are on the order of 1 TB. Currently, the largest data marts are on the order of 1 TB, while data warehouses are approaching 10 TB in size. It is not unreasonable to predict that these numbers will grow 10x or even 100x by the end of the decade.

Applications operating on very large data sets, such as database systems, perform best when as much of their data is in memory as possible rather than on disk. The access times for today's DRAM systems are roughly 100 times faster than today's fastest disk drives. Databases are on the forefront of applications that have an almost unlimited need for memory because of their large working sets. To efficiently address large amounts of memory beyond the 4 GB limits of today's 32-bit x86 systems, these RDBMS require 64-bit processors to support physical and virtual addressing powerfully and efficiently.

> **"Over the last ten years, the applications we've put on PCs have grown. They've grown in size and computational demands. And 32-bits of address space just isn't enough anymore. The size of databases has grown to the point where we just can't get the performance out of the 32-bit address space that we need to get to continue to support these applications. Over the past few years, we've added a few features to extend the life of the 32-bit system, but it's not enough, and we need to move to 64-bits to continue to support these large databases and high-end desktop applications."**
>
> **Dave Cutler**
> **Senior Distinguished Engineer, Microsoft**

At present, high-end RISC servers routinely ship with 8, 16, 32, or even 100 GB of physical memory. Low-end servers today are going out with 512 MB or 1 GB. As prices for DRAM continue to drop, it has become very cost effective to build systems with greater than 4 GB of memory. With the price of memory currently below $250 for 1 GB of ECC DDR memory, 8 GB costs less than $2,000. Because of the continuously declining costs of memory, the expectation is for low-end systems to be shipping with 2, 4, and 8 GB over the next two or three years. These systems must be able to address physical memory beyond the 4 GB limit. Current x86 processors support accessing memory beyond this limit through extensions such as Addressing Windowing Extensions (AWE) and Physical Address Extension (PAE), but they are non-portable, cumbersome to program, difficult to optimize performance, and are limited to accessing at most 64 GB of memory. A flat 64-bit address model is necessary to overcome these limits to gain maximum performance. The AMD Opteron processor can address up to 1 TB (1000 GB) of physical memory and 256 TB of virtual address space. This will enable systems based on the AMD Opteron to efficiently use memory well beyond the 4 GB limits of today's 32-bit x86 processors.

"All I know is that people keep using lots of memory, and every time memory gets cheaper, they use more of it. And so, as the memory prices come down, people put more in systems. Once they put it there, they want to be able to address it and use it. And that is where the 64-bit equation comes in."

Nathan Brookwood
Principal, Insight64

The AMD Opteron Processor

The AMD Opteron processor is AMD's eighth-generation processor designed for the server market. AMD Opteron ushers in a complete set of microarchitectural advances while fundamentally altering the next-generation of x86-based system architecture. The result is not only a generational leap in processor performance, but also a tremendous level of scalability in a system's delivered performance. The microarchitecture offers native support for 32-bit x86 software and features support for 64-bit computing with the AMD64 ISA. This combination of scalable performance for today's x86 software while enabling the use of the AMD64 ISA will lead to a new level of delivered performance, making AMD Opteron a true eighth-generation x86 processor.

The AMD64 Instruction Set Architecture

The primary factor that has enabled the unprecedented success of the x86 processor is that successive processor generations have retained software compatibility with previous generations. This allows users to maintain their investment in existing software while upgrading applications to the next generation at their own pace. AMD's 64-bit strategy is one that allows the latest in processor innovation to be brought to the existing installed base of 32-bit applications and operating systems while establishing an installed base of systems that are 64-bit capable.

AMD's strategy of extending the x86 architecture for 64-bit computing is a straightforward alternative to total conversion using incompatible instruction sets. The AMD Opteron processor implements the AMD64 ISA which is designed to permit platform suppliers, developers, and corporate IT departments to transition to 64-bit environments gradually while continuing to run 32-bit applications at full processor speed. By providing a smoother migration to 64-bit computing, AMD's strategy is designed to save its customers billions of dollars in software re-development and deployment costs.

Integrated Memory Controller

As processor capabilities have advanced, one of the greatest limitations on performance gain has become the system architecture's ability to deliver sufficient memory bandwidth to the processor core while reducing memory access latency. The AMD Opteron microarchitecture directly addresses this bottleneck by integrating a memory controller into the processor, revolutionizing the way an x86-based processor accesses main memory. The result is both greatly increased bandwidth available to the processor as well as reduced memory latency.

The AMD Opteron microarchitecture incorporates a 128-bit DDR DRAM controller interface as shown in Figure 1. The controller is designed to support PC1600, PC2100, and PC2700 DDR memory. This translates into available bandwidth to the processor of up to 5.3GB/s with PC2700 memory. The direct interface significantly reduces the memory latency seen by the processor and latency should continue to drop as the processor frequency scale.

The high bandwidth and low latency of the AMD Opteron processor's integrated memory controller has significant benefits when running a database system. Client requests to OLTP databases are typically random. In a banking system, requests from ATMs spread all over the world may be initiated simultaneously. The requests will be for unrelated data and will translate to requests from disparate parts of the database. Ideally, all of the index tables and the data itself will be present in memory. Accessing the data with as low a latency as possible will

Figure 1 – AMD Opteron Processor Block Diagram

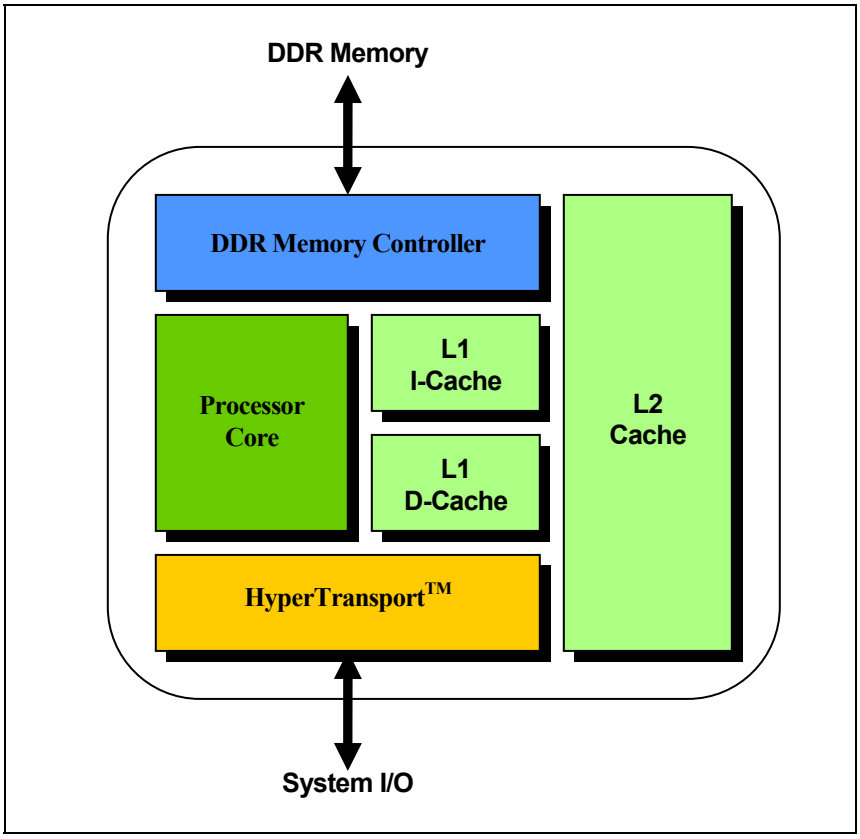

minimize the time it takes the RDBMS to turn a request into data to be returned to the client. The ability to process transactions quickly is the primary measure of the performance of a database system.

HyperTransport™ Technology Interface

The integrated memory controller has an even more dramatic effect in AMD Opteron processor-based multiprocessing systems. HyperTransport™ Technology results in an outstanding advance in x86 system architecture scalability by enabling glueless multiprocessing where the available memory bandwidth to the system scales with the number of processors. In Figure 2, an example of a four-processor server system is shown. In this configuration, the system is able to support up to 32 DIMMs capable of delivering up to an extraordinary 21.3GB/s of aggregate memory bandwidth to the system with PC2700 memory. Using 1 Gigabyte DIMMs, this four-processor system can be configured with as much as 32 GB of memory.

Figure 2 - Example Four-Processor Server System

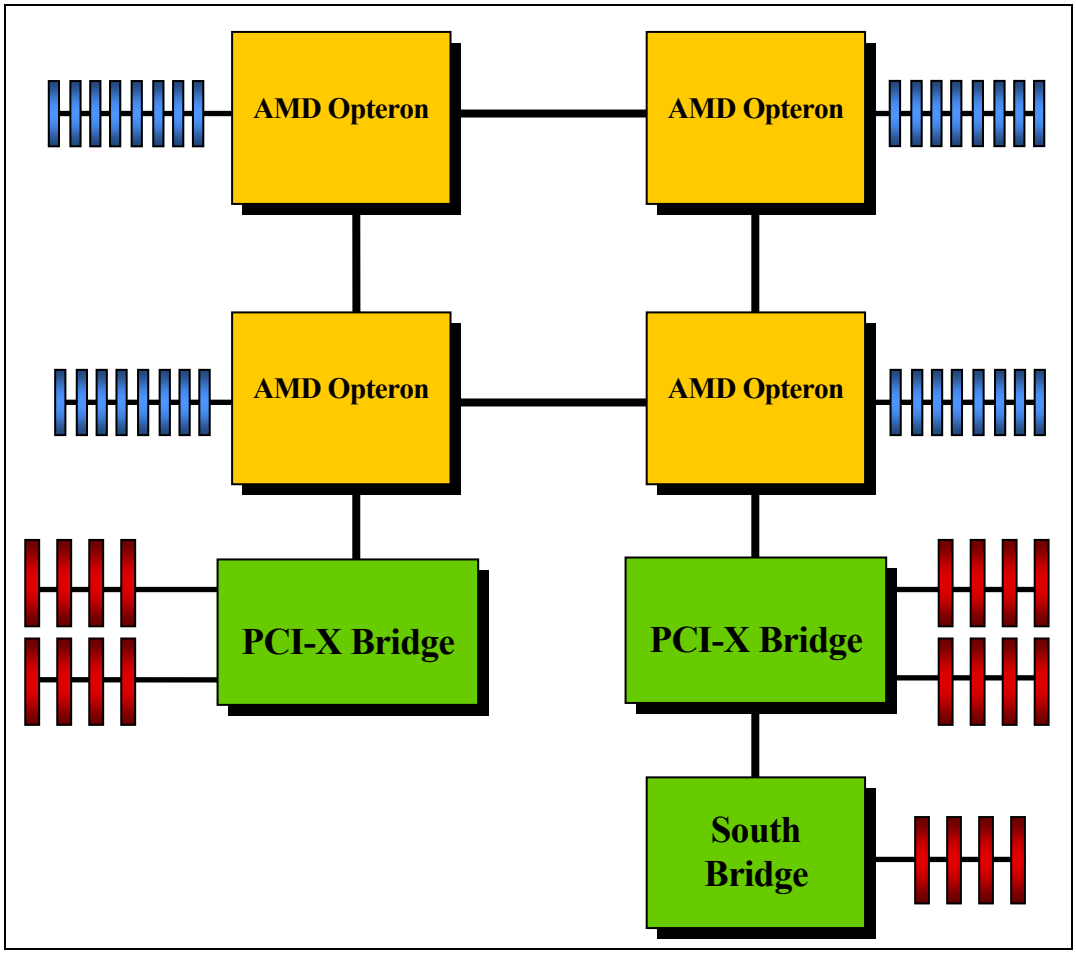

HyperTransport Technology links are built into AMD64 processors to provide scalable bandwidth interconnect links between processors and I/O subsystems. The links within the AMD Opteron processor are configured for 16-

bit communication in each direction at operating frequencies of 1600MT/s (megatransfers per second), for a maximum bandwidth of 3.2GB/s in each direction. Because of HyperTransport Technology high bandwidth and low latency, accesses to memory attached to a remote processor are significantly lower than in today's shared-bus SMP systems.

HyperTransport enables lower-cost systems by allowing multiprocessor systems to be built without expensive north bridge devices. This may further enable the trend from expensive scale-up systems to clusters of 2-way and 4-way systems. Not only are these clustered systems less expensive to build, but they enable the flexibility, performance, and reliability of clustered systems.

"With DB2 enablement on AMD64 technology, SuSE Linux shows enterprise-class application support on the industry-standard x86 platform. By combining the benefits of 64-bit computing and the x86 architecture, systems based on the upcoming AMD Opteron processors represent a next-generation server platform for Linux and will help accelerate the adoption of Linux in the enterprise."

Boris Nalbach
CTO, SuSE Linux

Mixed 32-bit and 64-bit Execution

One of the most compelling features of the AMD64 ISA is its support of 32-bit and 64-bit applications coexisting under a 64-bit operating system. More importantly, both 32-bit and 64-bit applications will run at full processor speed offering huge advantages for software vendors, OEMs, and end users. This coexistence allows software vendors to initially port only those applications that can take advantage of 64-bit computing and leave others as 32-bit applications. Porting less code translates shorter porting time, less testing, and less performance tuning leading to lower costs and faster time to market for the software developer and the OEM. End users do not have to replace all their applications, only those that can take advantage of 64-bit computing. For example, consider a system that is running both a database and an application server. The database application can be ported to 64-bit and the application server can remain as 32-bit. This allows end users to preserve their investment in 32-bit software while initially upgrading only those applications that can best take advantage of 64-bit computing. Additional applications can be upgraded to 64-bit as they are needed and become available.

Figure 3 shows a block diagram of the DB2 architecture running on a separate machine from an IBM WebSphere® client. The 64-bit DB2 database is running one machine while the WebSphere Application Server product (WebSphere) is running on a separate machine. The two applications are communicating over a network via TCP/IP.

Figure 3 - DB2 and WebSphere Running on Separate Systems

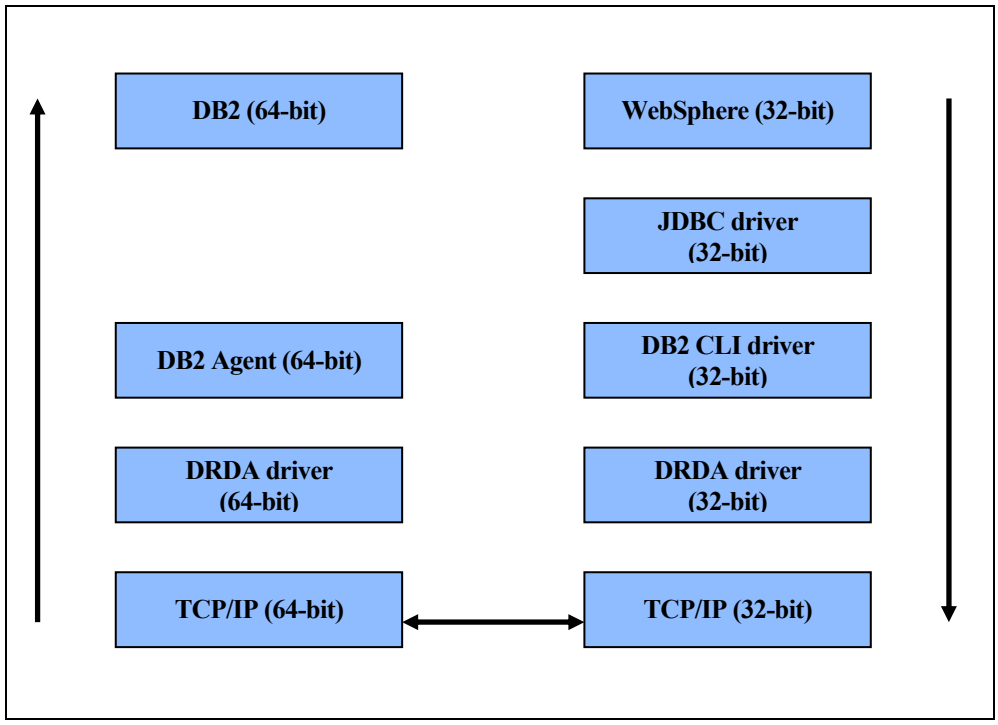

Figure 4 shows a mixed execution environment with 64-bit DB2 running on the same machine as the 32-bit WebSphere. DB2 has been ported to become a native 64-bit application while WebSphere is running unmodified as a 32-application. This preserves the user's investment in 32-bit WebSphere while the DB2 takes advantage of 64-bit addressing. Both applications will run at full speed. With the current server consolidation trends, IT departments want to do more with fewer systems. The ability to execute both 32-bit and 64-bit applications at full processor speed on a single system enables the consolidation of applications from two or more systems on to a single system.

"IBM is committed to support the efforts of its partners to enhance the scalability and reliability of Linux for the enterprise. The combination of DB2 and AMD Opteron processors is designed to enable customers to maximize their 32-bit applications and take advantage of native access to DB2's 64-bit environment."

Lauren Flaherty
Vice-president of marketing, IBM Data Management Solutions

Figure 4 - DB2 and WebSphere Running on an AMD Opteron Processor-based System

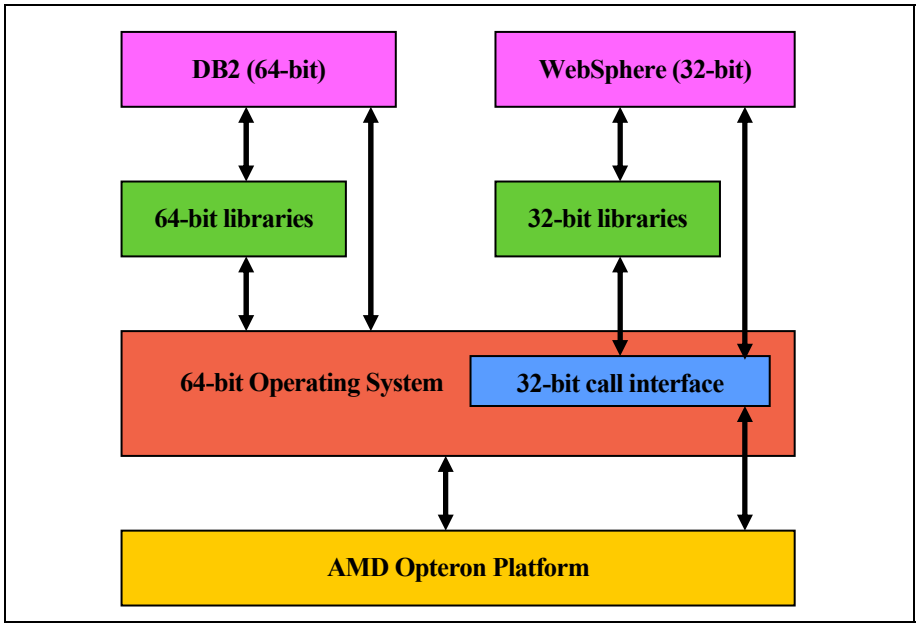

Ease of Porting to AMD64

IBM's DB2 source code has been ported to both 32-bit and 64-bit processor architectures. Before the port to AMD Opteron was started, the DB2 code was already "64-bit clean", meaning that all the code and data structures were written to support 64-bit computing.

The IBM engineering team was able to port several million lines of DB2 source code to 64-bit SuSE Linux in a matter of two days. This compares to a previous best-case time of several weeks to port DB2 , with some ports to other architectures taking many months. One of the most amazing aspects to consider about the port to AMD Opteron is that the only source code changes that were necessary were to a few makefiles – the application source code was not required to change at all.

Because the AMD64 ISA is an extension to the x86 instruction set, it allows software engineers to leverage all the years of knowledge and experience with x86. It allows them to take advantage of tools and techniques that they have learned over years of programming the x86 architecture. The unprecedented ease of porting is a testament to the robustness of the AMD64 ISA, software development tools, SuSE Linux, and AMD Opteron processor.

"The speed and ease of enabling DB2 to run on a 64-bit AMD Opteron processor-based system is a testament to the evolutionary nature of the AMD64 architecture. DB2 customers will have the power of 64-bit database computing in an x86 environment through the AMD64 architecture."

Richard Heye
Vice President, Platform Engineering and Infrastructure, CPG
AMD

Summary

The AMD Opteron processor offers the advantages of 64-bit processing while maintaining the advantages of the industry-standard x86 architecture. The architectural advances of the processor such as the AMD64 Instruction Set Architecture, the integrated memory controller, and HyperTransport technology make it ideally suited as a database server. IBM was able to port several million lines of DB2 source code in two days - faster than they have ever before.

AMD Overview

AMD is a global supplier of integrated circuits for the personal and networked computer and communications markets with manufacturing facilities in the United States, Europe, and Asia. AMD produces microprocessors, Flash memory devices, and support circuitry for communications and networking applications. Founded in 1969 and based in Sunnyvale, California, AMD had revenues of approximately $2.7 billion in 2002. (NYSE: AMD).