Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

# WHITE PAPER

# The AMD Opteron Processor: A New Alternative for Technical Computing

Sponsored by: AMD

Earl Joseph, Ph. D.
Christopher G. Willard, Ph.D.
Nicholas J. Kaufmann

November 2003

## A BRIEF HISTORY OF HIGH-PERFORMANCE COMPUTING PROCESSOR CHOICE

The high-performance computing (HPC) community has traditionally helped drive computer architecture and processor innovation, beginning with computing systems developed in the 1940s, to address defense and scientific applications. The initial stage of HPC processor use was necessarily based on proprietary processors, beginning with the first scientific processor designs and progressing through vector architectures and specialized processors for Massively Parallel Computers. This approach produced the fastest and most effective computing systems of their times and developed architectural concepts that were widely incorporated into subsequent processor architectures. These architectural designs had strong requirements for matching the processor speed to the memory system speed. Proprietary technical computing processors and memory systems became more specialized and expensive to produce over time, which limited their market growth.

The second stage in HPC processor use came with the introduction of reduced instruction set (RISC) computers. RISCs leveraged manufacturing techniques that allowed for single-chip processors and, subsequently, Moore's Law improvements in speed over time. Although these processors did not outperform proprietary architectures, they provided acceptable performance at a significantly lower price point and continual advances in both absolute and price performance. Computer system vendors developed these processors. Although the processors were made available to business partners, the vast majority of chips were delivered in the vendors' own systems. RISC processors combined with scalable system architectures allowed companies to amortize processor development cost across a broad range of technical and commercial markets, which in turn changed the economics of technical computing and expanded access to supercomputing technologies to a larger number of technical organizations. RISC-based systems have come to dominate the HPC market in terms of both systems revenue and shipments; we expect this technology to be a major player for the foreseeable future.

The third stage in the progression of HPC processor use began with work on commodity-clustered systems in the mid- to late 1990s. This cluster strategy assembles distributed-memory computing systems from "commercial off-the-shelf" (COTS) components. These systems were based on standard third-party processors. We use the term *standard* or volume computer processors to describe CPUs designed and manufactured by chip companies such as AMD that sell these components to a large community of computer systems vendors for integration into a wide range of computer products. It is interesting to note that standard processors were originally developed for desktop systems and have moved upmarket as a result of increasingly sophisticated designs and Moore's Law increases in absolute performance. (Moore's Law describes the doubling of computer chip density, and thus capability, approximately every 18 to 24 months). More recently, standard processors designed to meet the requirements of higher-end servers have come onto the market.

This white paper considers the AMD Opteron™ processor's entry into the high-end standard processor arena in light of requirements for high-performance systems markets. The paper provides a thumbnail sketch of the processor and highlights several features of the processor that may be of particular interest to HPC users. It presents a brief case study of the ASCI Red Storm development project — a high-end specialized system based on the Opteron processors. Finally, it provides IDC's analysis of the opportunity for standard processors in the HPC market as well as the challenges and opportunities for AMD in the market.

---

## AMD OPTERON OVERVIEW: BRINGING 64-BIT COMPUTING TO THE X86 WORLD

### ***OPTERON FEATURES FOR TECHNICAL COMPUTING***

The Opteron processor implements the x86 instruction set with a 64-bit memory space. The processor runs 32-bit x86 programs in native mode without changes and provides a 64-bit mode for running 64-bit applications. The processor provides program-controlled execution in either 32-bit or 64-bit mode. 32-bit applications can run on top of a 64-bit operating system. The compatibility for 32-bit jobs is in the microcode where there is only a small penalty for the conversion to a fixed-length instruction set. The improvements in silicon include out-of-order execution, enhanced branch prediction, improved translation look-aside buffer (TLB), and speculative execution. AMD increased the parallelism by first converting to RISC-like micro-operations (OPs), with deeper scheduling queues, out-of-order issue, and improved branch prediction.

The AMD Opteron processor is designed to provide balance between memory bandwidth and processor speed. For years the HPC industry has seen substantial increases in processor peak performance speeds without the same rate of improvement in memory bandwidth capability. The AMD Opteron architecture improves memory and system bandwidth in two ways:

- AMD integrated the memory controller into the AMD Opteron processor. This lowers latency and increases the effective bandwidth to memory. It also eliminates the front side bus. The memory controller no longer runs at the bus speed but at the processor clock speed. Furthermore, as processors are added to a system memory, controllers are also added so the bandwidth to memory scales directly with additional processors. Each processor has an independent connection to its local memory. So the memory in the system is physically divided between the processors. When a processor wants to address memory in another processor, the request goes through the HyperTransport™ logic and is handled by the local memory controller. There is no single central memory for which the processors contend.

- The on-chip interconnect interface uses AMD's HyperTransport technology. This high-bandwidth interconnect facilitates communications between processors and I/O. The HyperTransport processing is on the chip but in a separate unit from the CPU. Interconnect processing occurs without interrupting the memory controller or the CPU.

The Opteron architecture, with the memory controller and HyperTransport on the chip, delivers improved memory bandwidth scaling. The incorporation of the memory and I/O control on each processor chip removes the need for additional chips, circuits, and wires that add to latency and reduce the potential bandwidth. AMD Opteron latencies are presented in Table 1. Compared to AMD's Athlon™ processor, the Opteron has up to 3.9 times the bandwidth per processor and up to 38% lower latency in a multiprocessor configuration. For example, in a two-processor system, the Athlon MP's peak bandwidth is 2.7GBps or 1.35GBps per processor; in the Opteron system, peak bandwidth is 5.3GBps per processor. The Opteron's page-hit latency at 100ns is 38% faster than the Athlon MP's at 160ns.

## TABLE 1

AMD OPTERON MEMORY LATENCIES

| LOCATION OF REQUIRED DATA | PAGE-HIT LATENCY (ns) | PAGE-MISS LATENCY (ns) |
|---|---|---|
| Local memory | 65 | 95 |
| Adjacent processor | 100 | 120 |
| Two "hops" to processor | 140 | 160 |

Source: IDC, 2003

### MEMORY SPACE

AMD Opteron processor–based systems can support up to 1TB of physical address space. The memory controller can drive up to eight DIMMS and address up to 32GB of physical memory per processor. If we assume 2GB DIMMS, then it would be possible to build a 64-processor cache coherent Non-Uniform Memory Architecture (ccNUMA) with 16GB of attached memory per processor for a total of 1TB of memory. In such a system, processors would be arranged in a ring topology.

### COMPILERS

Several 64-bit compilers are available for AMD Opteron–based systems, including:

- GNU Compiler Collection (gcc) versions 3.2 and 3.3. The GNU 3.3 compiler is optimized for the Opteron processor.

- PGI Workstation 5.1 includes optimized Fortran 77/90 and C/C++ compilers. The compilers provide 32-bit and native 64-bit code generation, OpenMP support, and debugging support. 32-bit and 64-bit Linux and 32-bit Windows are available, with support for 64-bit Windows planned.

- Absoft is offering Fortran toolsets for AMD 64-bit technology for Windows and Fortran.

- NAGware f95 Compiler by NAG supports Fortran 95, high-performance Fortran, and legacy Fortran 77 code.

- AMD Core Math Library (ACML) includes Basic Linear Algebra Subroutines (BLAS) levels 1, 2, and 3; Fast Fourier Transforms (FFTs); and Linear Algebra Package (LAPACK). These AMD64 Instruction Set routines include Fortran and C interfaces and the ability to address single, double, single-complex, and double-complex data types. In addition an optimized version of the ATLAS BLAS library is available.

### DEVELOPMENT TOOLS

Additional tools are available for AMD Opteron processor, including:

- Blackdown Java Platform 2 version 1.4.2
- GNU C library (glibc)
- PERL
- 64Express code migration tool
- RSA encryption libraries
- MPICH (Open Source)
- Scyld Beowulf
- Debuggers, including Distributed Debugging Tool, a 64-bit graphical debugger; GNU debugger (GNB); SoftICE device driver debugger; and TotalView
- Code analysts tools, including AMD CodeAnalyst, OProfile, and Vampir/Vampirtrace

## AMD'S APPROACH TO PROVIDING APPLICATION COMPATIBILITY

### *32-BIT MODE WITHIN A 64-BIT PROCESSOR*

A typical strategy for making significant changes to established technologies is to provide dual modes of operation for the new version of the computer system technology: a previous-generation compatibility mode and a new-generation mode. This is done to make it easier for customers to move to the new computer. The AMD Opteron processor accomplished this by providing the following three support modes:

- **Legacy mode —** Preserves binary compatibility with both 32-bit applications and 32-bit operating systems. Programs previously written for 32-bit x86 processors need not be recompiled in order to run.
- **Compatibility mode —** Supports binary compatibility with existing 32-bit applications running under a 64-bit operating system.
- **New-generation mode —** Supports 64-bit applications and 64-bit operating systems. Applications must be recompiled to take advantage of 64-bit system features.

The Opteron processor can run a mix of 32-bit and 64-bit programs and switch modes with each change in program. An important feature for independent software vendors (ISVs) and developers is that AMD "didn't change anything" in the 32-bit mode. For example, in a Linux cluster environment, users and application developers can simply run 32-bit applications as they are without reprogramming or recompiling. At the same time, users can run both 32-bit and 64-bit applications on the same processor under the same OS.

The move to 64-bit applications is due to the need for larger memory, larger file sizes, and, in some cases, increased precision.

The challenges with moving to 64-bit include:

- Availability of open source software, ISV applications, and development tools
- Cost of reprogramming, major recompiles, and, in some cases, the complete rewriting of application codes
- The long and complex recertification process provides an advantage for many production-class codes to first run in 32-bit mode and slowly migrate to 64-bit processing as confidence grows

## CASE STUDY: ASCI RED STORM AT SANDIA NATIONAL LABS

### *THE PROJECT*

Sandia National Labs has contracted with Cray Computer Inc. to develop a large AMD Opteron–based supercomputing system for the lab. The project is designed to deliver a computer within 24 months leveraging Cray's MPP experience from the T3E product. Sandia will use the computer code named Red Storm to run its current set of applications in the ASCI nuclear stock pile stewardship program, including structural mechanical codes, computational fluid dynamics (CFD), climate models, and many classified applications. Sandia will provide access to the new computer to Lawrence Livermore and Los Alamos National Labs, the five ASCI university alliance partners, and Oak Ridge National Lab.

### *SANDIA PERSPECTIVE*

Sandia selected Opteron because of its price/performance, full compatibility with its IA-32 applications, and the open HyperTransport interface. Sandia's current large system is called ASCI Red built by Intel. The new ASCI Red Storm plans to be 7 to 10 times more powerful than the current system. Sandia will keep the older system for quite a while and move the largest, most challenging problems to the new AMD system. The new system architecture was designed by Sandia employees.

### *CRAY PERSPECTIVE*

Cray selected the AMD Opteron processor for building the ASCI Red Storm project for a number of reasons:

- It fits well with Sandia's current set of applications and won't require rewriting codes. It will initially run most of its applications in 32-bit mode, then move to 64-bit over time to take advantage of new features and capabilities.
- The open nature of the HyperTransport interconnect architecture provides an interconnect fabric for a 10,000-plus processor system.
- The processor provides an upward migration path to 64-bit requirements.
- Sandia's benchmark results showed that the AMD Opteron could deliver good performance on its codes.

Cray expects that the AMD Opteron processor will also provide good price/performance and high reliability.

### CRAY'S ROLE

AMD manufactures the processors. Cray's role will be to design and manufacture the node boards, memory and system interconnect, I/O, cabinets, cooling, and the overall operating system. Cray will also provide the system services and programming environments in Linux.

### SANDIA'S ROLE

Sandia will provide the computational portion of the operating system. Externally the computer will look like a Unix/Linux computer. Sandia will port any required ISV codes to 64-bit versions.

In the late 1980s Sandia developed a lightweight computational operating system to run on its nCube MPP system. Over the years, this OS has been continually enhanced and renamed. Intel developed a version called Cougar for the ASCI Red computer. Sandia's approach is to divide OS functions between compute and service kernels. The majority of nodes run the lightweight compute kernel, which is designed to minimize system overhead operations and thus maximize compute performance. A relatively few nodes manage OS services and run the service kernel. The small size and simplicity of the OS kernels have the added advantage of increased system reliability.

### SANDIA'S VIEW OF SYSTEM BALANCE

Processor speeds have advanced so quickly over the past decade that they have outstripped the ability of the memory system to provide data. Memory bandwidth and latency have become the dominating bottlenecks in supercomputers. In Sandia's words: "A good computer needs at least one byte of memory addressing per FLOPS; we expect ASCI Red Storm to deliver two per FLOPS."

Reliability is a key design feature, so Sandia designed the system with a real-time monitoring of all system components that have a separate RAS network within the computer. Sandia also felt that complexity had to be kept to a minimum, so there is only a single global operating system, and it is designed around uniprocessor nodes to avoid the complexity of symmetrical multiprocessors (SMPs).

### THE COMPUTER

The computer system will have over 10,000 AMD Opteron processors for computations and over 1,000 for handling system services. There can be 1–8GB of memory per processor in the design. The computer will require 124 cabinets and 3–4MW of power. It will be air cooled.

The computer interconnect will be based on a 3D mesh. Each processor will be a node with its own portion of memory. The first-phase computer is targeted to be installed in the second quarter of 2004, with the full system installed in late 2004.

Compilers include Fortran, C, and C++. Additional software includes Argonne MPI, Etnus TotalView, SCI libraries, and the PGI compilers.

Total disk bandwidth will be 100GBps, and the network has to provide 50GBps real bandwidth to the applications. The bandwidth of ASCI Red Storm is 25 times greater than that of the current supercomputer, ASCI Red.

### FUTURE PLANS

Cray plans to productize the project and sell smaller versions. Sandia has the rights to productize the computer if Cray does not bring it to market.

## IDC ANALYSIS: DEFINING SYSTEM BALANCE

### *BALANCED COMPUTERS AND DELIVERED PERFORMANCE*

The key to delivering high computational throughput in a computing system is a function of balance between processor speed, I/O throughput, memory bandwidth, and latency. HPC as an industry has lost ground in the 1990s, with processor speedups outstripping memory subsystem capabilities.

For more than a decade, processor performance has improved at a very high rate, doubling every 18 to 24 months. At the same time, the speed of memory systems has improved by a small percentage per year. So there is a growing difference in the speed that processors can calculate compared to the speed that data can be moved into and out of the processors. As processor performance outstrips memory performance, more applications become memory bound; thus, increasing processor speeds have little impact on overall application performance. Therefore, system performance is often limited by the balance between processor performance and memory bandwidth.

### *PERFORMANCE ATTRIBUTES*

There are many attributes that contribute to the performance of a computer. For different HPC users, the importance of these attributes varies based on the algorithmic structure of the problems being solved, the size, mix of the workloads generated by multiple users, as well as the urgency of solving key problems. Three major attributes of a computer's capability are:

- **Processor performance.** The desired capability is the speed at which the processors could generate results if they were kept fully occupied with work. Typically peak processor performance measurements only take into consideration the processor's clock rate, the number of instructions or other work that theoretically could be accomplished in each clock cycle, and the number of processors in the system. They ignore any bottlenecks created by the memory subsystem.

- **Memory system capability.** Actual performance is dependent on how much data can be moved into and out of the processors in a given time period. Key attributes include cache size, bandwidth, and latency; local memory size, bandwidth, and latency; and main memory size, bandwidth, and latency.

- **Scaling capability.** Systems that scale to larger sizes often require additional interconnect hardware and software to support communications between and coordination of multiple nodes. These extra components increase both delivered or available performance and price but do not affect peak performance rating. Factors affecting performance include number of nodes that can be configured, node size (processors and memory per node), total number of processors, total memory size, bandwidth between nodes (access rates to remote memory), and latency over the interconnect fabric. Larger nodes are more expensive to build but are generally easier to use, place less stress on the interconnect, and often provide higher performance across a broader set of problems.

Two primary measurements are used to describe the capability of a given memory system:

- **Memory bandwidth.** The amount of data that the memory system can handle in a given period of time; often expressed as gigabytes per second or billions of bytes per second.

- **Memory latency.** The time it takes to start getting data to and from memory, expressed as either the time to move the first data element from the time it is requested to the time it is received. Sometimes it is also expressed as the number of clock ticks required to receive a data request.

Figure 1 presents memory bandwidth on the left axis compared to memory latency on the right axis for communicating between an increasing number of processors in the system. The actual scale changes over time, so the scales in the charts are meant to show relative conditions. The best case is when bandwidth (the left axis) is always very high and latency (the right axis) is always low.

**FIGURE 1**

HIGHEST-PERFORMANCE DESIGN APPROACH

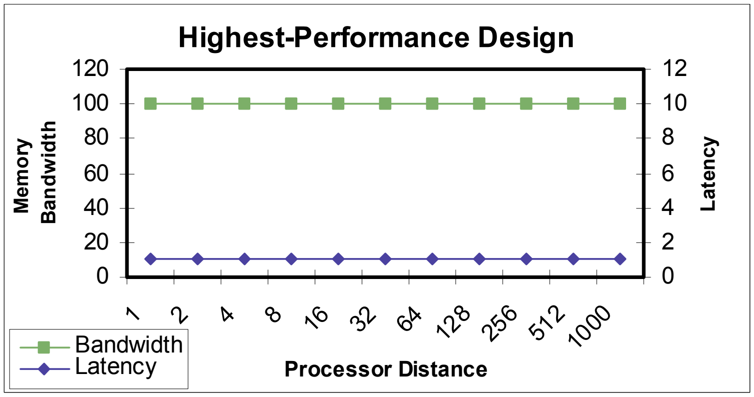

Note: Bandwidth and latency values are relative performance numbers.

Source: IDC, 2003

One of the easiest approaches is to design a computer with very high bandwidth and low latency that matches the improvements in processors, but this results in a computer that is very costly (tens of millions of dollars).

A second approach is to construct technical computers using single processor compute nodes to reduce the high costs of adding memory bandwidth subsystems. With a single processor node, the design only needs to connect one processor to memory, whereas in multiprocessor designs, all of the processors need to be connected to the memory system, and logic must be added to allow the sharing of the memory between the processors. Traditional MPPs and clusters of PCs are examples of this approach (see Figure 2). This is typically the lowest-cost approach and can provide good solutions for running many small capacity jobs and/or for jobs that are parallel. Due to the lower costs, this is the highest-growth approach.

## FIGURE 2

LOWEST-COST DESIGN APPROACH

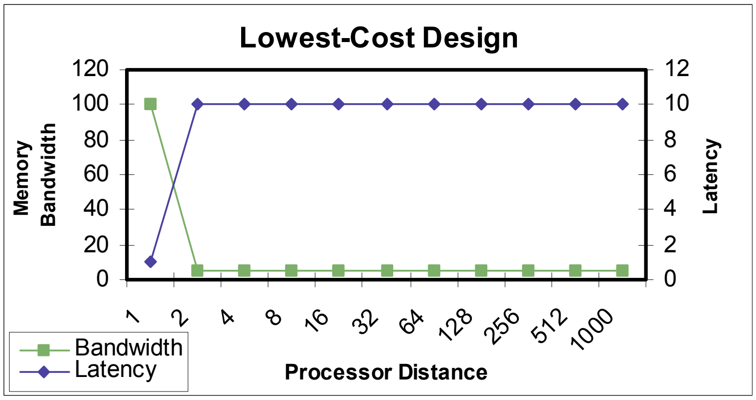

Note: Bandwidth and latency values are relative performance numbers.

Source: IDC, 2003

A third approach that provides good performance on small and medium-sized jobs and scales without a very expensive memory subsystem is clustering SMPs (see Figure 3). An SMP provides a single high-performance, flat-memory subsystem that supports a large number of processors (typically 32 or 64 processors). SMP clusters work well when most of the capacity jobs fit within one SMP node.

## FIGURE 3

THE SMP DESIGN APPROACH

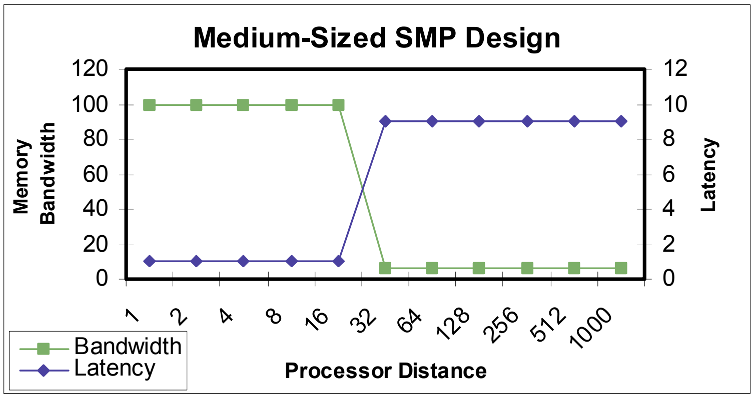

Note: Bandwidth and latency values are relative performance numbers.

Source: IDC, 2003

The NUMA strategy works to balance both cost and memory system performance by building servers with a memory subsystem that delivers strong memory capability to a small node, and then instead of dropping memory capabilities drastically, the memory subsystem capability is reduced at a moderate rate as the system size grows (see Figure 4). These computers fit well for HPC users who have a broad mix of job sizes and for large jobs that aren't highly parallel. Other types of NUMA approaches use lower-performance interconnects. In this example, we are referring to NUMA design with higher-performance interconnects.

## FIGURE 4

THE NUMA DESIGN APPROACH

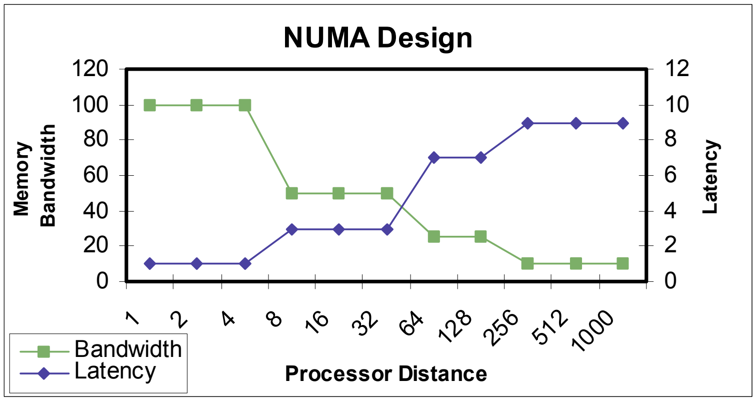

Note: Bandwidth and latency values are relative performance numbers.

Source: IDC, 2003

## AMD'S APPROACH TO PROVIDING BALANCE

The AMD Opteron was designed with the goal of improving the balance between bandwidth and processor speeds. The AMD Opteron architecture supports balanced memory performance by integrating the memory controllers and HyperTransport technology into the processor chip.

### *HYPERTRANSPORT TECHNOLOGY*

HyperTransport technology was designed to be a higher-speed, lower-latency, chip-to-chip link for connecting peripheral circuits and processors on a system board.

AMD developed its HyperTransport technology as a high-speed, low pin count, asynchronous, chip-to-chip, board-level interconnect (clock forwarded, point-to-point bus). AMD has designed its HyperTransport technology to be an open standard through the HyperTransport Consortium. Communications delays caused by interchip latencies are eliminated by placing the memory controller and HyperTransport technology on the Opteron chip, and the performance of the communications system is coupled, at least in part, to the clock speed of the processor.

The Opteron's HyperTransport implementation has three 16-bit links, which run at up to 1.6 gigatransfers (800MHz DDR). Each Opteron has three links with 6.4GBps of bandwidth each.

## FUTURE OUTLOOK — STANDARD PROCESSORS AND CLUSTER COMPUTING

Over the past decade, the HPC market has seen a dramatic technology shift away from systems based on proprietary processors and toward systems based on commodity RISC processors and standard or volume processors. This latter category of CPUs is designed and manufactured by chip companies, such as AMD, and incorporated into computer systems by computer manufacturers. Figure 5 presents the HPC market systems revenue by commodity RISC, proprietary, and standard processors — that is, the value of computer systems categorized by the type of processor each system employs. In 1993, commodity RISC processors accounted for about $198 million or 13% of the total market revenue; by 2000, this number had increased to over $4.6 billion or 92% of revenue. More important, commodity RISC–based systems accounted for virtually all of the growth in the market. In 2001, standard processor–based systems began to appear in the market in noticeable numbers and, in 2002, accounted for $605 million or nearly 13% of revenue.

### FIGURE 5

SALES OF CUSTOMER COMPUTERS COMPARED WITH SALES OF COMMODITY COMPUTERS

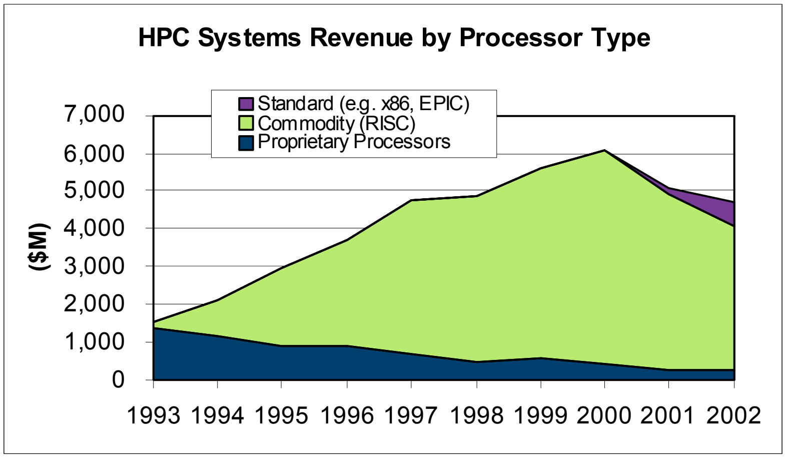

Source: IDC, 2003

In absolute terms, standard processor–based systems accounted for $132 million in HPC systems sales in 2001 and $605 million in 2002. It is important to note that about 93% of this latter value was associated with x86 architecture processors. Although it is too early to make specific predictions, IDC believes this data indicates a real change in the processor mix, and we expect standard processor–based systems to continue to gain market share.

We believe the recent success of standard processors is driven by a combination of absolute processor power, system price/performance, and cluster technology. Of particular interest is the success of 32-bit processors in technical markets. Technical computing has strong requirements for 64-bit precision, high floating-point computation speeds, and large memory address spaces. We believe that 32-bit processors have succeeded to date based on a combination of price/performance (i.e., even systems running slower double precision calculations are still cost effective), growth of 32-bit applications in the biosciences market, and the distributed memory nature of clustered systems. We expect the entry of 64-bit processors, such as Opteron, to expand the proportion of the overall HPC market that can be effectively addressed by standard processor–based systems.

### ***COMMODITY PROCESSORS AND CLUSTERS***

Industry-standard processors have succeeded in the technical computing markets in large part through their use in clustered systems configurations. If we consider only clusters sold by traditional HPC vendor, we note that in 2002 clustered systems generated about $584 million in HPC market revenue; this value represents a 225% increase over cluster sales of $180 million reported in 2001 (see Figure 6). Systems based on standard processors dominated this market with a 73% share of revenue in 2001 and a 99% share of revenue in 2002. (It is important to note that these shares can show significant yearly fluctuations based on a few sales of very high-end systems. Even so, we believe standard processor systems hold and will continue to hold the lion's share of the market.)

IDC estimates that the $584 million in cluster revenue for 2002 represents approximately 15,000 system shipments, with these systems containing about 170,000 processors. We expect the technical cluster market to grow to about $1.5 billion and 34,000 system shipments by 2006, with processor shipments to outpace both revenue and shipment growth. Finally, we should note that the above analysis represents only systems sold through traditional HPC channels and identified as clusters by vendors. Another "dark cluster" market exists for systems sold through nontraditional channels and/or assembled by end users. This market is potentially as large as the traditional HPC vendor market for clusters. However, we expect this "do-it-yourself" market to diminish somewhat as more vendor-based alternatives become available.

## FIGURE 6

HPC CLUSTER REVENUE BY PROCESSOR TYPE

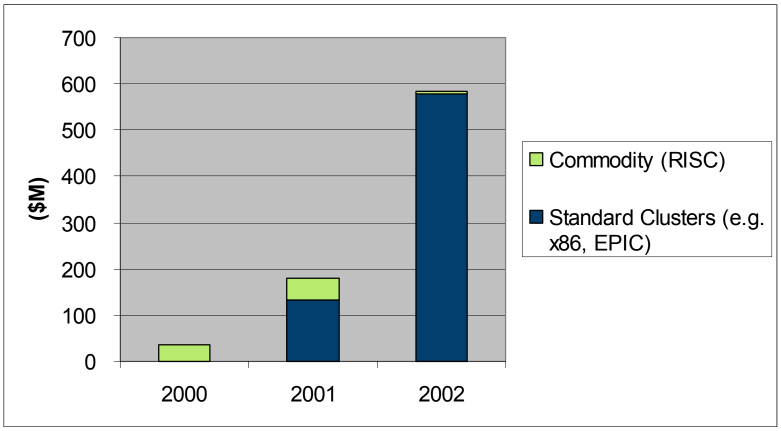

Source: IDC, 2003

The move to clusters demonstrates two important characteristics of technical markets:

- **Technical innovation.** Scientists and engineers continually increase requirements for computer system capability. These requirements lead to a willingness to work with new concepts and technologies. Subsequently, technical markets are often the entry point for new processors and architectural strategies, and thus they can provide a good entry opportunity for Opteron processor–based systems.
- **Technical advancement.** As initial users of technologies, HPC organizations have traditionally added value to the new systems through software development and integrations. This willingness to support new technologies provides an opportunity for vendors, such as AMD, to accelerate product development and market entry.

Cluster market growth could be accelerated by 64-bit, standard high-performance processors. This segment represents the best initial opportunity for Opteron because of the high penetration of x86-based clusters. Follow-on opportunities in the technical market would include:

- **Product expansion.** AMD can work to leverage success in the cluster market to grow the Opteron ecosystem and expand the use of the processor into a large set of products and markets.
- **Opteron-based products.** Options are available to build products for the broader technical computing market based on the Red Storm architecture. In addition, AMD can work to develop partnerships with additional HPC system vendors to field Opteron-based systems beyond clusters.

## CHALLENGES

Technical computing markets will quickly adopt new technologies that can demonstrate real capability advantages to end users as well as staying power in the market and functional growth. We believe that in meeting these requirements, standard processor technologies must address several specific challenges for entering the technical market:

- **Develop an OEM base.** Users evaluate the degree of standardization of a technology based on how many ISVs supply and support that technology. New standard processors in the market will need to be incorporated into a number of different OEM vendor products.
- **Develop a technology ecosystem.** In addition to OEM support, new technologies need to develop a support infrastructure or ecosystem of organizations with products and/or expertise associated with the technology. These organizations include ISVs for applications and middleware, systems integrators, and system support/training organizations.
- **Provide ongoing technical support.** Standard processor vendors must provide continued support of compilers, tools, and libraries. This support can be provided either directly by the vendor or through partnerships.
- **Maintain system performance and balance.** Technical computing customers demand products that provide high performance on their applications and tend to look at the performance attributes first in making buying decisions. Vendors must demonstrate that the products not only can deliver strong performance at any given time but also can be expected to grow performance over time with minimal sacrifice of system balance.

## MEETING THE CHALLENGES

The design of the Opteron processor, with its unique ability to run both 32-bit and 64-bit workloads simultaneously plus the HyperTransport interconnect technology, positions AMD to address many of the above challenges. By making the HyperTransport technology an open standard, AMD will help motivate additional market acceptance and OEMs. The Sandia-Cray partnership puts AMD on the radar screen for many HPC buyers. It gives AMD initial credibility and will lead many user organizations to at least take a look at AMD before buying their next server.

# CONCLUSION

The Opteron processor and its associated technologies offer a refreshing new alternative to the HPC server market. HPC customers are always looking for new technologies and alternative system designs as a way to create competitive advantage through HPC. Given the economic pressures of the past few years, buyers are less willing to make investments in costly porting and rewriting of software. The ability to run their current software and applications without changes and then enhance applications over time provides a strong opportunity for AMD.

## COPYRIGHT NOTICE