IT TAKES AN **ECOSYSTEM** TO MAKE A DUAL-CORE PLATFORM

Part II in a Series of AMD White Papers on 64-bit Computing

The sweep of history includes a never-ending series of battles between conflicting ideas: Athenian democracy versus Spartan monarchical oligarchy; the geocentric universe of Ptolemy and Aristotle versus the heliocentric universe of Copernicus; the Church of Rome versus the Protestant Reformation; and capitalism versus communism, to name a few.

Similarly, the development of the Information Technology (IT) market has been marked by ideological as well as technological conflicts: mainframes versus minicomputers; vertical versus horizontal integration; proprietary versus open architectures. One of today's battles revolves around the best way to develop and deliver contemporary computing platforms. Some suppliers believe they can do it all themselves. They develop and deliver processors, chipsets, communications solutions, and even barebones systems. These suppliers have even been known to use their intellectual property portfolios as a barrier to limit third-party competition with their own products. AMD takes a different approach. They believe customers are best served when they have an opportunity to deploy best-of-breed technologies in their IT platforms. They realize that they cannot do it all themselves. AMD partners with the best chipset, I/O controller, and software developers. In short, they believe it takes an ecosystem to make the best dual-core platforms, and AMD has developed just such an ecosystem.

In this white paper, we provide an overview of the breadth and depth of the AMD64 ecosystem from both a hardware and software perspective, and highlight areas where IT practitioners should pay special attention as they plan their system infrastructure. We touch on the thorny software licensing issues created by the move to dual-core and multi-core processors, and discuss AMD's position with regard to these issues. But before we start the tour of this ecosystem, we want to share with you our views regarding the significance of the industry's new focus on dual-core and multi-core processors, and what this shift means for buyers and users of computing systems.

Why the shift to dual-core?

Since the invention of the microprocessor in 1971, CPU designers have relied primarily on increases in clock frequency to drive advances in processor performance. Recently that approach has begun to lose its effectiveness, since increasing clock frequency forces a chip to use more power, which in turn makes it harder (and more expensive) to cool. Today, most processor architects believe there are easier and more cost-effective ways to improve system performance. The shift from 130nm to 90nm process geometries doubled the transistor budget available to chip designers and created the opportunity to put two processor cores on a piece of silicon no larger than a single-core chip produced with a 130nm process. Dual-core technology

can improve the performance of applications by almost 100 percent or more* compared with single-core technology, while traditional frequency scaling offers improvements of 10 to 20 percent at best.
Dual-core (and eventually multi-core) technology should be viewed primarily as a design technique to increase system performance, rather than as a way to convert dual-processor systems into less expensive uniprocessor configurations. The history of the computer industry provides ample evidence that workloads expand to absorb capacity, and dual-core technology is unlikely to change this phenomenon. The use of dual-core processors in DP and MP blade, rack, and tower configurations allows systems to handle bigger workloads (the same benefit that frequency increases have traditionally provided), potentially reducing the number of systems required to handle the aggregated load. This, in turn, can lead to reductions in rack space, floor space, input power, and HVAC expenses. It may also result in lower software licensing fees, depending on the software vendor's licensing models and the nature of the workload, as we will discuss later in this document. Dual-core technology has a place as well in desktop and mobile client systems. Performance-sensitive applications can benefit the most, as long as the application code has been designed to exploit multi-threading techniques. Even single-threaded applications can benefit, since operating system and housekeeping functions can run on one core, freeing the other to focus on application execution. Dual-core clients often "feel" more responsive to the end-user, since the operating system has more flexibility in scheduling resources and is less likely to "lock up" when a misbehaving application seizes control of one of the cores and refuses to let go for an extended period.

Dual-core technology can improve the performance of applications by almost 100 percent or more* compared with single-core technology, while traditional frequency scaling offers improvements of 10 to 20 percent at best.

Chipsets form the basis of AMD's hardware ecosystem

In the AMD64 world, chipsets play the key role of connecting processors to the subsystems and buses that handle disk, network, video, and other I/O-related activities. AMD's chipset partners include ATI, Broadcom's ServerWorks, nVidia, SiS, Uli, and VIA. These companies offer a range of chipsets that enable everything from low-end, value-oriented platforms with integrated graphics to high-end, eight-way SMP configurations for mission-critical deployments. Chipsets from these suppliers incorporate the latest technologies, including PCI-Express and SATA hard drive controllers with RAID capabilities.

In the past, a move to "the latest technology" often implied a need for updated device drivers, a problem for IT managers who want to maintain a common system image for all the client systems deployed over the course of a year. Some IT groups solve this problem by deferring their adoption of new technologies until their suppliers no longer manufacture the systems they sought to acquire. Others expend resources to update system images as the underlying hardware evolves during the year. AMD's new Commercial Stable Image Platform (AMD CSIP) addresses this problem by ensuring that the device drivers their key partners release before September 1st of a given year will work with the CSIP-enabled devices they release over the following fifteen months. Unlike some image stability programs that force customers to adopt a single supplier's chipsets and wireless devices, AMD's program gives customers a choice. They can select systems based on ATI or nVidia core logic, and Atheros or Broadcom wireless facilities. Whichever combination they select, customers can be confident that their subsequent purchases will operate properly with the drivers they install at the start of the stabilization period.

*http://www.amd.com/us-en/Processors/ProductInformation/0,,30_118_8796_8800~97024,00.html

Motherboards facilitate system design for OEMs and system integrators

Over 20 manufacturers of motherboards, including all of the "top-tier" and "second-tier" suppliers, offer desktop boards that support AMD64 processors. The world has come a long way since the 1999 COMDEX show, where many of these same suppliers requested NDAs before they would talk (behind closed doors) about their AMD Athlon™ processor-based boards. AMD64 desktop motherboards are available in microATX, ATX, and extended ATX formats. Seven of the leading suppliers that focus on the server market (Arima, Asustek, Gigabyte, IWILL, MSI, Supermicro, and Tyan) support the AMD Opteron™ processor. All provide support for two-way AMD Opteron processor-based configurations, some provide support for four-way, and two (IWILL and Tyan) even offer board sets that accommodate up to eight AMD Opteron processors in SMP configurations.

Barebones systems simplify the delivery of AMD technology

System integrators and sophisticated end-users can purchase AMD Opteron processor-based systems from a variety of suppliers. Several leading motherboard providers, including Supermicro and Tyan, sell barebones systems based on their boards. Contract manufacturers like Sanmina-SCI offer two- and four-way systems based on their own motherboards (that are not sold on a boards-only basis).

AMD's Direct Connect Architecture gives developers a few new options

Most I/O adapters used in personal computers and servers are attached to the processor via standard buses like PCI, PCI-X, PCI-Express, and AGP. Such devices will work with systems based on competitors' and AMD's own chips, although they may work a little better on AMD64-based systems due to AMD's superior Direct Connect Architecture. But a few designers have discovered that they can improve their products' performance by connecting their devices directly to one of the AMD Opteron processor's HyperTransport™ technology (HT) channels. PathScale, the first company to release a product based on this concept, claims its InfiniPath HTX™ adapter delivers the lowest latency MPI interconnect available today (1.29 microseconds) for high-performance cluster configurations. The PathScale product essentially provides a direct bridge between an HT link and an InfiniBand switch fabric, and works with off-the-shelf InfiniBand 4X switching devices. You should see additional HT-enabled devices from other suppliers in the future.

Software for dual-core 64-bit computer systems

The Linux community was quick to exploit AMD64 technology; Novell (then SUSE) released a version of SUSE Linux adapted for AMD64 to coincide with the AMD Opteron processor launch followed by other Linux distributors. Sun released Solaris 10, its first version with AMD64 support, early in 2005, and Microsoft released 64-bit versions of Windows XP Professional and Windows Server 2003 x64 Edition a few months later, shortly after AMD Opteron celebrated its second birthday. In all, over 300 ISVs are now delivering hundreds of software packages for AMD64.

All major database engines have been certified on AMD64 with packages that run in the 64-bit extended mode.

The availability of these 64-bit operating systems paves the way for the emergence of 64-bit database environments. Oracle 9*i*, 9*i* RAC, 10*g*, and 10*g* RAC, IBM DB2 and Informix, CA Ingres, and Sybase have already been ported and released on AMD64 platforms. Microsoft's SQL Server 2005, the 64-bit adaptation of its popular SQL Server 2000, was released in November 2005. All major database engines have been certified on AMD64 with packages that run in the 64-bit extended mode. The 64-bit addressing capabilities of AMD Opteron™ processors allow database management software to utilize vast amounts of DRAM memory as a cache for frequently accessed records. Since DRAM accesses take microseconds, while disk accesses typically require tens of milliseconds, this caching technique dramatically improves throughput and responsiveness for database accesses, and often means you need fewer systems (and fewer database software licenses) to handle a given workload.

Compared with the massive effort required to establish the AMD64 ecosystem, it required virtually no incremental effort to create an ecosystem for dual-core processors. Software optimized for dual-processor systems runs without change on dual-core configurations. In fact, unless a program inspects some relatively obscure bits in a status register, it cannot distinguish between a dual-processor/single-core and single-processor/dual-core arrangement. The performance benefit that any specific application gains from a move to dual-core systems depends on a number of factors that we will detail below, but generally overall system performance will be enhanced by such a move, since dual-core platforms double the execution resources available to the scheduling software within an operating system.

AMD suggests that users characterize their prospective workloads with regard to two key parameters: (1) compute intensity versus memory and I/O intensity, and (2) single-threaded versus multi-threaded applications. In general, compute-intensive workloads gain more from a shift to dual-core configurations than do I/O—or memory-intensive loads. Such a result should come as no surprise, since Dual-Core AMD Opteron processors provide twice the compute cycle bandwidth but essentially the same memory and I/O bandwidth as single-core processors. The ability of any particular application to benefit from the increased computational bandwidth depends on its internal structure. Multi-threaded applications typically benefit more than single-threaded versions, since their application code can run in parallel on both cores. Single-threaded applications benefit less, since all application code runs on a single core, but they still can gain from the ability to run other applications and OS code in the second core. We have reviewed a variety of server and workstation workloads from the perspective of compute intensity and multi-threading. Figures on the following page illusrate the relationship between these workloads and dual-core technology.

Workstation/Scientific Workload Postioning

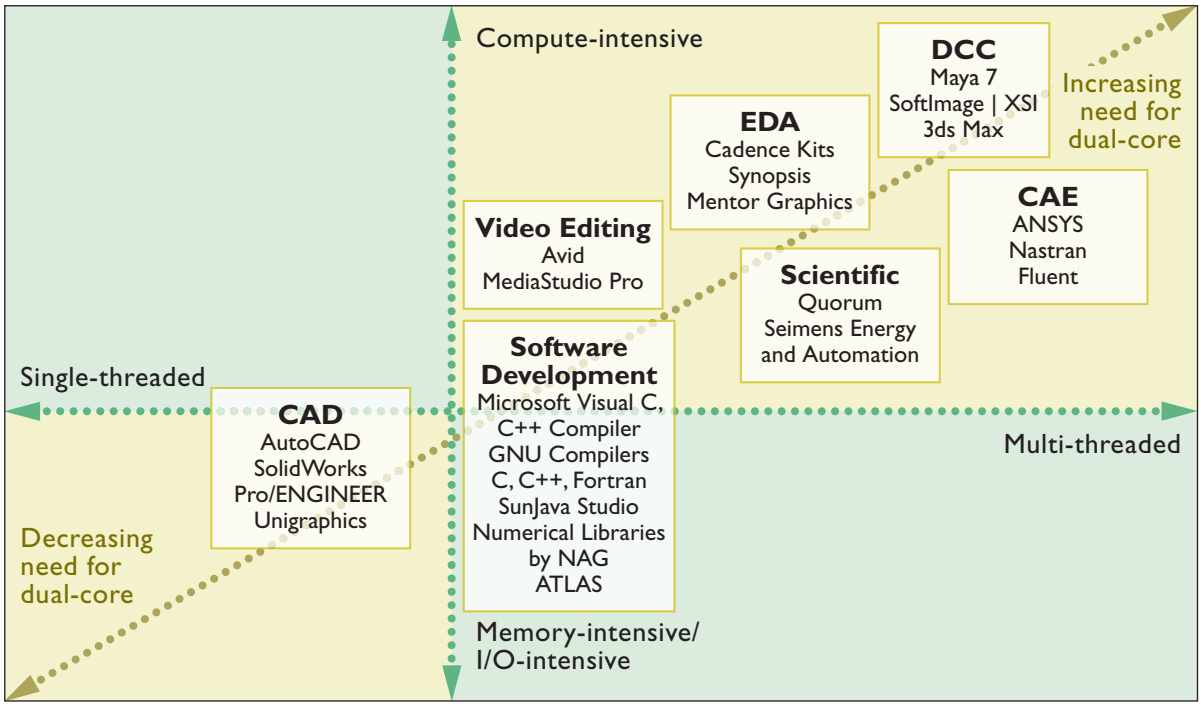

Server Workload Positioning

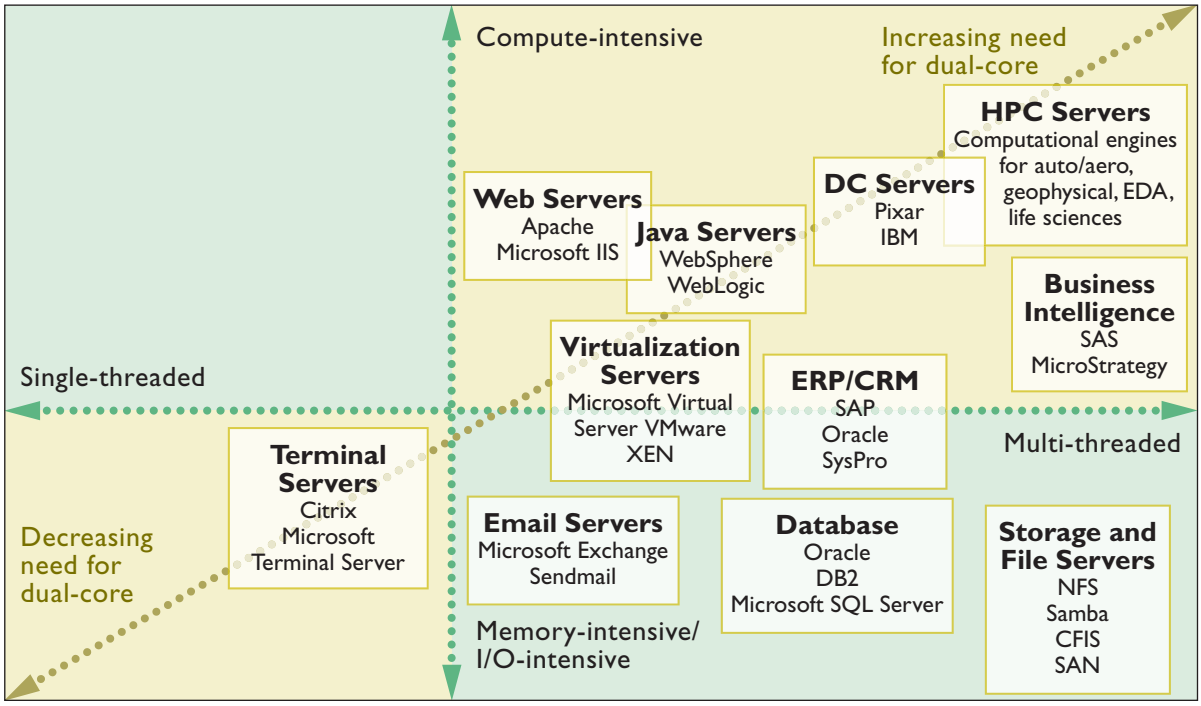

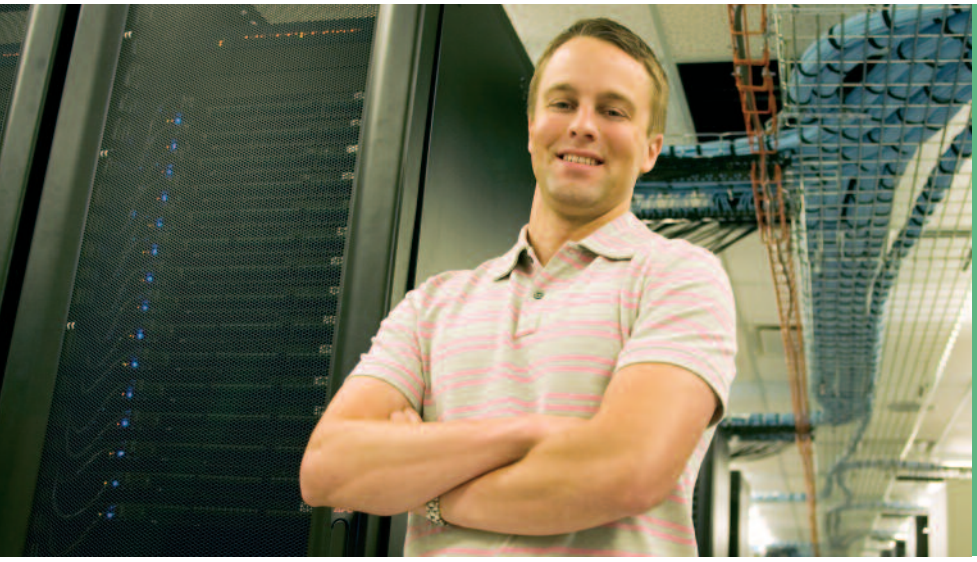

As the first processor supplier to deliver 64-bit extended x86 systems, AMD's platforms have acted as magnets for software developers in need of the highest performing x86 systems.

AMD's software ecosystem continues to grow

As the first processor supplier to deliver 64-bit extended x86 systems, AMD's platforms have acted as magnets for software developers in need of the highest performing x86 systems. There are thousands of packages certified for AMD64. These include:

- Operating systems, including Linux, Solaris, Unix, and Windows.
- Hundreds of development tools, including compilers, debuggers, Integrated Development Environments (IDEs), JVMs/JIT (Java), math libraries, profilers and performance analyzers, and source code control systems.
- A variety of database engines, including all of the popular models have been adapted for 64-bit operation, as we noted above.
- Hundreds of infrastructure applications, including Application Servers, Cluster Management, Desktop Productivity Software, Directory Services, Email/Groupware Servers, File and Print Services, Server-based Computing (SBC) Servers, Storage Management, System Management, Virtualization, and Web Servers.
- Packages that address a variety of vertical market applications in areas such as CAD, CAE, EDA, 3D scientific visualization, digital content creation, decision support, ERP, and CRM. These packages operate in both 64- and 32-bit modes.

The list of certified packages grows daily. You can get the latest updates by checking AMD's Web site at www.amd.com/amd64ecosystem.

Software licensing for dual-core 64-bit computer systems

Although users may install identical software packages on single-core and dual-core configurations, the license fees they must pay to use this software can vary significantly. Software suppliers tend to fall into one of two camps in this regard. Enlightened ISVs (including Microsoft and IBM) realize that dual-core technology is just another way to improve microprocessor performance. Since few have scaled their prices to reflect MPU performance in the past, they see no reason to charge more for the software that runs on a dual-core processor than they charged for software that ran on the single-core chip that preceded it. They ignore the number of cores on a chip when calculating their software license fees, and continue to count only the number of processor chips in the system.

Not all suppliers are so enlightened. Some software vendors take the position that an execution core is an execution core, regardless of whether it shares its silicon with other execution cores. These vendors argue that their software delivers twice the value when it runs on a dual-core chip, and therefore they charge twice as much for that software. AMD encourages all ISVs to price their software realistically with regard to the single-core/dual-core design philosophy.

Virtualization: the next frontier

The philosophy of "one application per server" has driven the growth of the x86-based server market over the past decade. This approach has led to data centers filled with racks filled with underutilized systems, each over-provisioned with regard to CPU, DRAM, and hard drive resources. It has now become clear that, in order to improve overall resource utilization, it will be necessary first to virtualize the workloads running on these underutilized systems, and then to consolidate these virtual environments onto a smaller number of physical systems. HP's "Adaptive Enterprise," IBM's "On Demand Environment," and Sun's "NI Provisioning" all require that machine resources be virtualized, so that they can be flexibly assigned and reassigned in response to external requirements. Mainframe architectures long ago added hardware that facilitates virtual machine environments, but x86-based systems are only now beginning to address these issues.

AMD has worked with the key developers of x86 virtual machine environments, including VMware, Microsoft, and XenSource, to ensure that their current software-based virtualization technologies are certified for operation on AMD Opteron™ and AMD Athlon™ 64 platforms. AMD has also worked collaboratively with developers in these organizations to design a set of AMD virtualization processor extensions, formerly codenamed "Pacifica," that are designed to reduce virtualization overhead and harden the technology. AMD has released the virtualization technology specifications to these developers, along with an emulator package that facilitates development and testing of software that supports them. AMD plans

AMD realizes that no vendor has a monopoly on good ideas, so they have recruited the best and the brightest hardware and software partners to help them deliver the best dual-core platforms you can buy today.

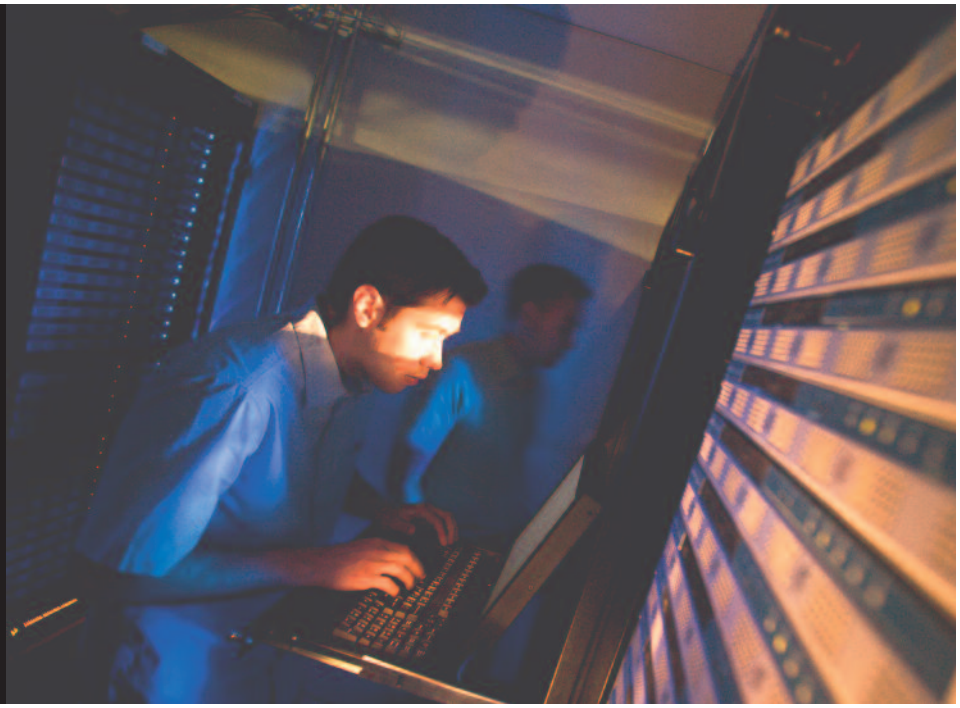

on including support for these extensions in future generations of AMD Athlon™ 64 and AMD Opteron™ processors. If you want to learn more about AMD's plans for virtualization, be sure to visit www.amd.com/virtualization.

Summary

We hope this document has convinced you that the ecosystem behind a processor matters almost as much as the chip itself. AMD is proud of the microprocessors they produce, but we understand that those processors require robust systems to deliver their full potential, and even then would be useless without the right software running on them. AMD realizes that no vendor has a monopoly on good ideas, so they have recruited the best and the brightest hardware and software partners to help them deliver the best dual-core platforms you can buy today. As a result, Dual-Core AMD Opteron processors give end-users more choice in the systems they can purchase, and give system partners more choice in the components they can use to assemble those systems.

Of course, AMD knows that it takes more than a solid architecture and a solid ecosystem to make an industry-leading multi-core processor. They started with a great architecture. Then they built an ecosystem to support that architecture. They equipped factories with the latest manufacturing processes and process management technology to stamp out vast quantities of processors based on this architecture.

If you have found this document to be informative, we hope you will eagerly await the next one in our series.

40379A