Insight
64

Industry-standard 64-Bit Microprocessors Will Permanently Alter the High End Systems Landscape

Nathan Brookwood
Principal Analyst
(Nathan@Insight64.com)

September 2002

In the three decades since Intel introduced the first microprocessor, advances in semiconductor manufacturing technology and processor architecture have dramatically increased the capabilities of these devices. These improvements have driven massive changes in the structure of the computer industry.

During their first decade, microprocessors enabled the creation of programmable calculators like the HP35 and its successors. During their second decade, they drove the growth of the personal computer market. By the end of their third decade, industry-standard microprocessors accounted for almost 90 percent of server sales, although these units were heavily skewed toward the low end of the market.

As the computer industry enters the fourth decade of the microprocessor era, suppliers have positioned their most advanced chips for an assault on the high end of the computer systems market. Although these more powerful systems represent only five percent of unit shipments, they account for over 65 percent of industry revenues, and thus form an attractive target.

Established system suppliers have responded to the emergence of these more powerful industry-standard processors in varied manners. Some, like HP, have moved aggressively to exploit the advantages of new industry-standard processor technology, such as the Itanium processor family. Others, like Sun, insist the old ways are best, and continue to pursue the strategies that drove their past success. Still others, like IBM, straddle the fence, dabbling with both old and new approaches.

In this white paper, Insight 64 assesses the state of current and future 64-bit microprocessor technology, and examines the pros and cons of industry-standard and proprietary microprocessors. We conclude that 64-bit industry-standard microprocessors offer overwhelming technological and economic advantages over the proprietary processors that dominate today’s high-end markets. Over the next five years, these advantages will drive a massive shift to systems based on industry-standard 64-bit processors.

I. Introduction

For millions of years, the waters of the mighty Colorado River flowed unimpeded from the Rocky Mountains to the Gulf of California, etching the Grand Canyon and many lesser canyons into the landscape of the Southwest. In years with unusually high precipitation, the river overflowed its banks, flooding adjacent lands. These periodic floods wreaked havoc in areas along the river that had developed into important agricultural centers by the early part of the 20th century. In 1930 the US government invested $49 million to control this flooding. Over a period of five years, construction crews poured 4.3 million cubic yards of concrete into an area known as Black Canyon and formed the Hoover Dam, standing 590 feet above the river's bed. The dam blocked the Colorado's flow, and over the next four years small towns named Rioville, Callville and Junctionville disappeared under 400 feet of water in the newly created Lake Mead. A repositioned Callville now hosts one of the largest inland marinas in the country. Hilltops 100 miles upstream of the dam now go by the names Little Boulder Island and Bighorn Island. The construction of a state-of-the-art dam on a scale beyond anything ever attempted forever changed the landscape in the southwestern United States, and repositioned the area for economic growth few envisioned at the time.

For over thirty years, an unending stream of improvements in semiconductor technology has allowed processor designers to dramatically increase the performance and capabilities of the products they create. Early 4-bit and 8-bit microprocessors lacked the power to handle general purpose computing, but enabled the emergence of sophisticated handheld calculators that wiped venerable slide rule producers like Keuffel & Esser off the map. 16-Bit processors like Intel's 8086 and Motorola's 68000 powered the desktop PCs and Macintosh computers that flourished in the 1980's. These systems delivered substantial operational and economic advantages over the time-shared minicomputers that had been used for interactive computing. Formerly successful minicomputer suppliers like Prime Computer and Wang Laboratories found themselves underwater, never to regain their former glory. Intel's 32-bit 486 and Pentium processors powered entry-level servers well matched to the needs of small workgroups. As the capabilities of these 32-bit chips grew, they assumed a larger role in the IT infrastructure. By 2002, 32-bit x86-based servers accounted for almost 90 percent of all server systems, and all the major server suppliers[1] included x86-based server systems in their product lines. Advanced processors forever changed the landscape of the low end systems market.

Although the market for systems priced below $25,000 has coalesced around the x86 architecture, the market for higher-priced servers remains fragmented. Most systems incorporate one of the RISC architectures that emerged in the 1980's, including HP's PA-RISC, Sun's SPARC, IBM's POWER, and HP's Alpha. A few trace their lineage as far back as the IBM System/360 architecture that first appeared in 1963. High end systems often include substantial amounts of physical DRAM memory, and thus were the first to add 64-bit features that facilitate memory access beyond four gigabytes. Servers often support thousands of on-line users in mission-critical applications, and thus their processors, along with other system components, incorporate a wide variety of features to enhance system reliability and data integrity. Those who market these systems have credibly argued that systems constructed with CPUs originally designed for desktop applications[2] lack the reliability and data integrity features needed for high end applications. But now, Intel has launched its second generation 64-bit Itanium® 2 processor, and AMD plans to launch its 64-bit Opteron™ processor early in 2003. Both suppliers claim their designs incorporate the reliability and data integrity features this segment of the market demands, along with the 64-bit addressing features needed for large memory system configurations. These new industry-standard 64-bit processors have to date gained an inconsequential share of the market. Insight 64 believes that by the end of the decade, these processors will dominate the high end systems market in much the same way x86-based designs dominate the low end. In the remainder of this document, we outline our basis for that belief and our view of the manner in which this transition will play out.

---

[1] Sun Microsystems, the final holdout, finally added the x86-based LX50 to its product line in August 2002.

[2] A somewhat transparent reference to x86 processors.

A Few Relevant Microprocessor Definitions

CISC (Complex Instruction Set Computing), the earliest processor design philosophy, originally aimed at minimizing program size and system memory requirements in an age when physical memory was both expensive and slow. Over the past decade, CISC designs have assimilated many architectural concepts that first appeared in RISC processors, and have closed the performance gap that once existed between CISC and RISC designs. Early CISC examples include IBM mainframes (originally the System/360, now known as the eServer zSeries), and DEC PDP 11 and VAX minicomputers. The CISC-based Intel 8086 eventually became the most prevalent processor architecture of all time.

RISC (Reduced Instruction Set Computing) emerged in the 1980's as an attempt to simplify processor design and improve performance. For many years, RISC processors outperformed their CISC counterparts by a wide margin, and RISC proponents argued that its performance advantage stemmed from the cleaner, more streamlined architectures of the RISC designs. More recently, both RISC and CISC designs have attained similar levels of performance, demonstrating that a processor's internal design has at least as much influence on its performance as its instruction set architecture. Today's most popular RISC processors include Sun's UltraSPARC series, HP's PA-RISC line, IBM's POWER processors, and the HP Alpha processor (originally developed by Digital Equipment, and subsequently developed by Compaq).

EPIC (Explicitly Parallel Instruction Computing), the newest processor design philosophy, comes from a joint HP/Intel development program in the 1990's. As its name suggests, the EPIC approach embeds in the instruction stream explicit information that tells the processor which program instructions may be executed in parallel. Prior approaches (i.e., CISC and RISC) require a processor to extract instruction-level parallelism within a program on the fly, a task that greatly complicates the overall CPU design. EPIC designs place the burden of extracting parallelism on the compilers that translate high level programming languages into low level machine instructions, a one-time task accomplished during the software development cycle. Intel's Itanium processor family is the first product line based on the EPIC concept.

Proprietary processors are sold only in conjunction with complete computer systems, and are not marketed as processors per se. Proprietary processors rarely appear on the merchant semiconductor market. Buyers normally associate proprietary processors with a single system supplier (e.g., SPARC with Sun, POWER with IBM, PA-RISC with HP), although other system suppliers occasionally offer use these processors in their lines (e.g., Fujitsu with SPARC, Bull with POWER and NEC with PA-RISC).

Industry-standard processors are sold on the merchant semiconductor market to original equipment manufacturers (OEMs) that incorporate these processors into systems sold to end users. Suppliers of industry-standard processors sometimes package their wares with other system building blocks (typically boards and chasses), and sell the resulting assemblies to OEM customers. Insight 64's definition of "industry-standard processor" relies on the processor supplier's position in the system component supply chain, rather than on the installed base or annual sales of the processor or processor architecture. Prominent general purpose industry-standard microprocessors include Intel's Pentium®, Celeron™, Xeon®, and Itanium® lines, the IBM/Motorola PowerPC®, the AMD Athlon™ and Duron™ series, and the planned AMD Opteron™ line.

32-Bit processors possess the ability to work with 32-bit virtual memory addresses. This feature makes it easy for software running on such processors to manipulate pointers within data structures as large as 4,294,967,296 bytes, commonly referred to as 4 gigabytes (GB). Conversely, such processors can deal with data structures larger than 4 GB only with great difficulty. Until recently, the cost of a 4 GB DRAM array ran into the tens of thousands of dollars, so this theoretical addressing constraint rarely limited actual programs. At current levels, a 4 GB bank of DRAM sells for a few thousand dollars, and this constraint has become a real problem for many applications.

64-Bit processors possess the ability to work with 64-bit virtual memory addresses. This feature makes it easy for software running on such processors to manipulate pointers within data structures as large as 18,446,744,073,709,551,616 bytes, commonly referred to as 16 exabytes (EB). Few (if any) programs in existence today need to address this much memory, but programs constrained by the 4 GB limit of 32-bit processors will find they can manipulate 8 GB, 128 GB, or even terabytes of data with ease following their adaptation for 64-bit operation. Programs designed to operate in a 32-bit environment must, by definition, be specifically adapted to exploit the extended addressing capabilities offered by 64-bit architectures. 32-Bit programs that run unmodified on 64-bit processors continue to behave like 32 bit programs, and face the same constraints as 32-bit programs running on 32-bit processors. No specific processor design philosophy (CISC, RISC or EPIC) has a monopoly with regard to 64-bit operation. Several 32-bit RISC architectures have been successfully extended to support 64-bit operations. The Alpha and Itanium architectures included 64 bit capabilities in their initial definitions. AMD's Opteron adds 64-bit extensions to the CISC-based 32-bit x86 instruction set architecture that has been in use since the 1986 launch of the Intel 386.

II. The Computer Systems Industry Continues Its Shift to a Horizontal Structure

In the early days of the computer industry, vertically integrated system suppliers had little choice but to design and manufacture virtually all the hardware and software components they used in their systems. In today's more mature industry, an evolving infrastructure allows suppliers to outsource many of the components needed to build desktops and servers. Adroit suppliers seize opportunities to increase their added value and differentiate their offerings through focused innovation in specific aspects of the system's design. These differentiated products almost always retain the ability to operate with industry-standard environments and applications packages.

Figure 1 illustrates the roles played by technology suppliers and system suppliers in the new horizontally structured systems business. Intel and AMD provide industry-standard processors, along with elements of the core logic needed to assemble systems around those processors. Microsoft and a bevy of Linux developers provide the operating systems, middleware and applications that run on these industry-standard platforms.

Figure 1
The Shifting Structure of the IT Systems Market

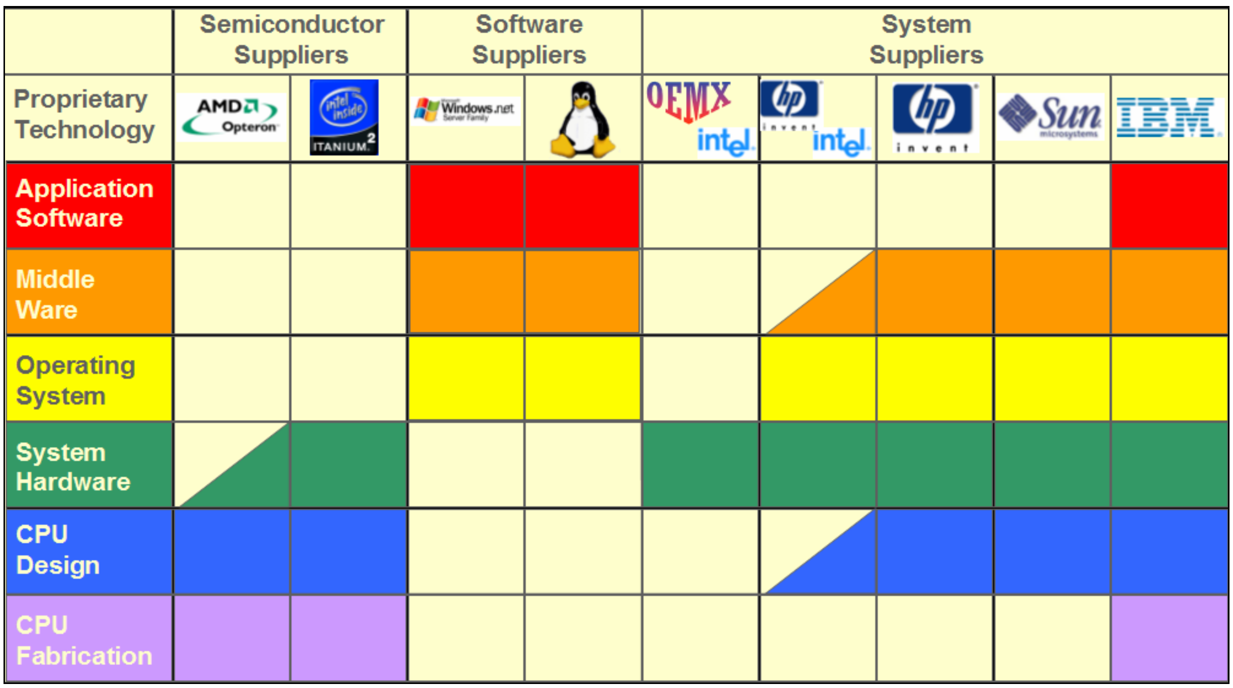

Source: Insight 64

The manner in which system suppliers augment the technologies they acquire from semiconductor and software suppliers varies considerably. The column labeled "OEMx/Intel" represents a new approach to high end system design, in which system OEMs apply their hardware design expertise to create differentiated platforms based on industry-standard processors. IBM's xArchitecture group follows this model, as do groups at Bull, Hitachi, NEC, and Unisys, among others. IBM's other hardware business units follow a more traditional model, and continue to rely on many internally developed technologies. They design and manufacture the processors used in IBM's proprietary pSeries, iSeries and zSeries offerings, wrap system hardware around those processors, develop the proprietary operating systems, middleware and applications that run on those platforms, and offer support for all aspects of the system hardware and software through the IBM Global Services organization.

Other RISC-based system suppliers pursue a fabless variation on the vertical integration theme. Hewlett-Packard's PA-RISC and Alpha-based product groups develop enhanced versions of their proprietary processors, along with the system platforms built around those platforms and much of the operating system software (HP-UX and TRU64 Unix) that runs on those platforms. HP relies on third parties to manufacture its PA-RISC and Alpha processors, although at an earlier time, these devices were manufactured in-house as well. Sun's approach mirrors that of HP's PA-RISC and Alpha groups.

Hewlett-Packard's position in the Itanium ecosystem differs from that of all other Intel OEMs in two key regards. As the co-developer of the Itanium architecture and several Itanium processors, HP contributes to the processor technology that it and other Intel OEMs use in their high end systems. HP also licenses its HP-UX operating system, the only "enterprise class" Unix implementation on Itanium, to other Itanium-based suppliers such as NEC and Hitachi that target the high end systems market. HP's involvement in the Itanium processor development program gives the company some unique advantages with regard to its Itanium-based systems, and allows the company to offer an almost seamless transition from its earlier PA-RISC product lines to the newer Itanium-based systems. For example, HP's next-generation PA-8800 RISC processor utilizes the same system bus structure and form factor as Intel's Itanium 2 processor. This allows HP to offer platforms that support PA-RISC or Itanium 2 processors, and are easily upgraded (in place) after a customer decides to migrate to the Itanium platform. HP's HP-UX transition strategy is equally accommodating, and allows customers to continue to use their PA-RISC binary software images in an emulation mode even after they have migrated to Itanium 2 processors.

III. The Coming Battle between Proprietary and Industry-standard 64-Bit Processors

The battle for control of the high end of the server market will be long and messy. The incumbent purveyors of platforms based on proprietary processor technology, including Sun and IBM, plan to defend their positions with all the technological and marketing skills they can muster. Their processor and system roadmaps contain a plethora of increasingly powerful systems as far as the eye can see. On the other side of the battle line, Intel, AMD and a wide range of system suppliers are gearing up for their assault on these established positions. Ultimately, as indicated in Figure 2, industry-standard systems will come to dominate shipments of high end systems, just as they dominate the low end today. Insight 64 believes the battle for the high end will be fought along the independent vectors of performance, scalability, system reliability, software availability and economics.

> **Performance** will clearly play a role in determining the winner, but it is neither the be all nor end all some might suggest. Sun's SPARC series has traditionally lagged both Alpha and PA-RISC on most computational benchmarks, but has managed nonetheless to gain the largest share of the market. To win the hearts, minds and pocketbooks of prospective system purchasers, any challenger must demonstrate it is at least as fast as the market leader, and it must convince buyers that it can maintain a competitive level of performance on an on-going basis. At the time of its launch during the summer of 2002, HP servers based on Intel's Itanium 2 processor registered the highest scores ever attained by a 64-bit processor on the SPEC CPU2000 benchmark suite, prompting one self-proclaimed long-time Itanium skeptic to issue a report titled "Itanium 2 Performance: *Wow!* "[3]
>
> Intel's Itanium line derives much of its performance edge through the use of advanced compilers that translate application code into machine instruction sequences. These compilers are just beginning to tap the performance potential of the EPIC architecture, with its unique combination of innovative features including predication, speculation and explicit parallelism. (Predication removes costly branches that impact performance while speculation allows the loading of instructions and data from main memory well before the processor needs them.) Enhancements to the suite of Itanium compilers will improve the performance of software running on installed systems, as well as on future platforms with faster processors. Intel's 64-bit roadmap, populated with exotic locales named Madison, Deerfield and Montecito, suggests the Itanium processor family will be anything but a one-trick pony with regard to performance.
>
> **Scalability** matters in the high end server market. HP's Superdome handles up to 64 processors and the SunFire 15K supports up to 106 (albeit with serious compromises regarding I/O capacity). Itanium 2 processors contain features like fast and flexible system buses and large on-die caches that enable OEMs to design highly scalable systems. The Itanium 2 processor's 128-bit, 6.4 GB/second bus gives OEMs the ability to develop a wide range of innovative core

[3] Jonathan Eunice, "Itanium 2 Performance: *Wow!*," Illuminata, August 27, 2002

logic solutions that best address the scaling requirements of the market segments they serve. OEMs targeting low volume, high end configurations with hundreds of processors can add the sophisticated (and expensive) switching logic that allows those processors to operate coherently and reliably. OEMs addressing high volume, low end configurations can create simpler designs that shorten their development cycles and reduce their costs. The Itanium 2's 3MB L3 cache, the largest ever incorporated on a microprocessor die, boosts performance and enhances scalability by offloading memory traffic from the system bus. The next generation Itanium processor due in 2003, code-named Madison, doubles the size of the L3 cache to 6MB, further enhancing the scalability of Itanium systems.

Much of the technology that makes a system "scalable" resides in the platform, rather than in the processor itself. Designers have yet to agree whether NUMA, SMP or clustered approaches yield the best performance for high end systems with a large number of processors. With so many Intel OEMs pursuing so many different Itanium-based system architectures, end users will soon have the ability to test and compare these varied approaches, and select the versions that best satisfy their unique set of requirements. Eventually, this Darwinian selection process may determine a "best" approach for high end system architectures. It will more likely demonstrate that each approach has pluses and minuses, and that customers prefer to select from a range of options.

Figure 2
Industry-standard 64-Bit Processors Will Dominate System Shipments by the End of the Decade

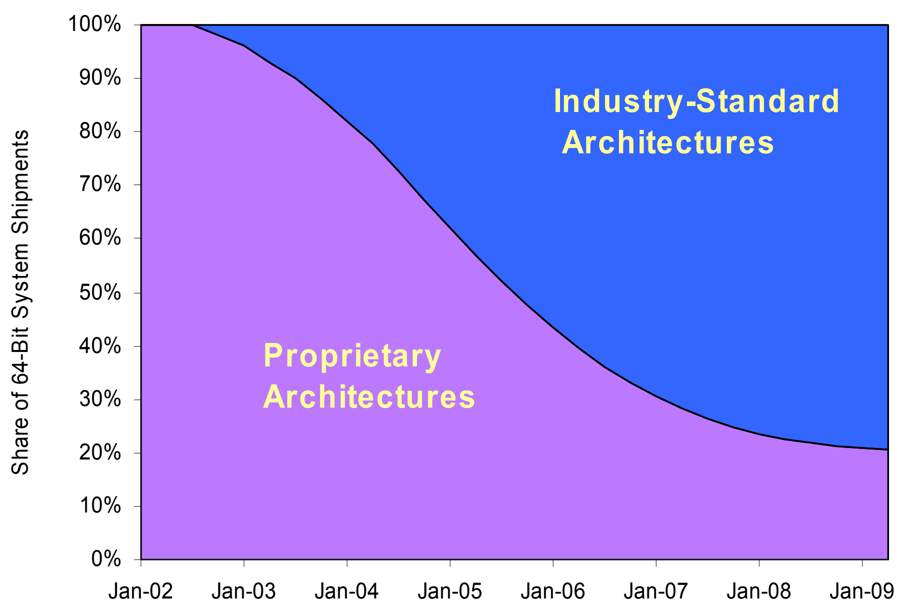

Source: Insight 64

Reliability is not an optional add-on for high end systems. Such machines are often deployed in mission-critical applications that operate on a 24x7 basis. Reliability and serviceability considerations must permeate all aspects of the system's design. This applies especially to the design of the microprocessor. The Itanium 2 processor's Machine Check Architecture (MCA) provides the features needed to monitor processor operation and ensure the reliability of both the CPU and the system. MCA monitors all critical paths and resources within the processor, correcting errors when possible, and disabling malfunctioning units before they can impact overall system operation. Without such features, it would be impossible to build highly reliable, enterprise-class systems. With these features, knowledgeable OEMs have the resources they need to build systems based on industry-standard processors that match the reliability of systems based on proprietary RISC processors.

Software plays a key role in determining the commercial success of any computer architecture. Without software, even the fastest systems become expensive room heaters with no useful function. In today's IT world, independent software vendors (ISVs) supply most software. These ISVs focus their development resources on the platforms they hope will maximize their software revenues. They sell some of their products to customers with previously-installed hardware; for products of this nature, as they make their investment decisions they pay special attention to the size of the installed base. Other software products tend to be purchased along with new hardware systems. Development decisions regarding products in this category must reflect the anticipated volume of platform sales over the life of the software package. In ordinary times, the hardware vendor with the largest installed base also has the highest volume of new system sales. This allows an ISV that looks only at the installed base to still be well positioned with regard to the stream of new system sales.

Insight 64 believes that the introduction of industry-standard 64-bit systems into a market previously served only by proprietary solutions creates a discontinuity that allows an industry-standard architecture with a small installed base to achieve a high volume of new sales. The discontinuity results from two key factors. First, as described below, the economics of the industry-standard business model enable compelling improvements in price and price/performance, as compared with proprietary systems. Second, the market share for an industry-standard platform equals *the sum of the shares* amassed by each system vendor offering versions of that platform. This aggregation of demand around a single architecture differs materially from the proprietary model, in which the market share for any given architecture approaches as a limit the market share of the supplier that promulgates that architecture. Consider, for example, that the total market share for the Alpha architecture is nothing more than the market share for HP's Alpha systems. By way of contrast, the market share for x86-based servers is the sum of HP's share of this segment, along with the shares of Dell, IBM, Fujitsu, Siemens, NEC, Bull, and many others. Astute ISVs recognized this unusual constellation of circumstances, and have begun to adapt their software for operation on Itanium systems in anticipation of the architecture's future gains in market share.

Industry-standard systems gain yet another software edge through their support of a diverse set of operating environments. Proprietary RISC platforms are typically offered with a proprietary operating system, often derived from Unix. In recent years, some vendors have added native Linux ports or Linux-like interfaces that run on top of their proprietary Unix adaptations. By way of contrast, 64-bit industry-standard servers offer 64-bit versions of Windows Advanced Server, the most popular server operating system, along with Linux. Hewlett-Packard adds a version of its enterprise-class HP-UX environment, and plans to add its Nonstop Kernel (NSK) and OpenVMS environments as well. Many Itanium system suppliers have licensed the HP-UX software and offer it with their systems. Hardware system suppliers gain the economies of scale that result when a single hardware architecture can be sold across multiple markets and with many software environments. Software developers gain the economies that stem from their ability to sell the same binary package to the customers of many hardware suppliers.

Economic issues will have a profound impact on the outcome of the battle between proprietary and industry-standard processors, and give the industry-standard designs a huge advantage. Beyond their obvious impact on a system supplier's R&D and manufacturing expenses, proprietary processor strategies increase the end-user customer's switching costs and limit the overall utility of systems based on proprietary designs. Let's examine each of these issues in detail:

> **High Development Expenses.** The design of microprocessors and operating systems is a complex, labor-intensive task. Sun boasts that it has 1400 engineers working on various implementations of its SPARC microprocessors. Insight 64 estimates this activity adds over $300 million per year to Sun's R&D expenses; Sun's figures exclude the engineers needed to design the computer systems that are based on these

processors. CPU designers need expensive CAD tools to do their jobs. Since it takes at least three years to develop advanced microprocessors, organizations need at least two teams, working on separate projects on staggered schedules, in order to introduce a new or enhanced processor every two years.

Proprietary processor strategies also entail substantial software investments. Operating systems must be ported; compilers and run-time environments must be developed or adapted; virtually every piece of code that runs on a proprietary system must be specially arranged for. This is not just a one-time expense or worry; each time an ISV upgrades a software package that runs on an industry-standard platform, special provisions must be made to move that new version onto the proprietary platform.

Ultimately, the buyers of proprietary systems must pay for all this development activity. The outlay to develop an industry-standard microprocessor is the same (or even a bit higher) than the cost to develop a proprietary processor, but these development costs can be spread over a vastly larger number of units. Consequently, the development expense per unit falls dramatically. Competitive market conditions ensure that much of this savings gets passed onto the end buyer of the system.

High per-unit manufacturing costs. The cost to manufacture any chip is a function of both its size and its production yield (i.e., the fraction of usable chips that results from any given production run of those chips). When a chip is manufactured in high volumes, the manufacturer can afford to run many sample production lots to understand production parameters and optimize yield. In a low-volume environment, fewer sample lots are used, and production yield remains low over the product's lifetime. This translates into a higher unit cost for each good microprocessor that results from a given production run.

High Switching Costs for End-Customers. As they deploy computer systems, customers invest in software specifically adapted to those systems. Software acquisition or development costs often dwarf the cost of the system hardware. Once an application is up and running and the staff is trained in its operation, customers understandably shun changes. A move to an incompatible hardware environment usually entails additional software and training expense, and greatly increases the risk something will go wrong along the way. In the end, high switching costs and increased risk induce customers to maintain old and inefficient systems instead of moving to newer, more efficient ones. Conversely, the users of systems based on industry-standard processors can migrate from one supplier’s systems to another’s, using the same operating system environment and binary software on both systems, and minimizing the cost of switching suppliers.

Reduced System Utility for End-Customers. A customer who wants to add a specific peripheral device or application program to a system may find that the device driver or application code has not been ported to the proprietary environment. This precludes the use of that peripheral or application on the proprietary system, and reduces the overall value of the system for that customer. (Economists refer to the advantages provided by a wide array of complementary products as "network effects." They are one of the main reasons customers prefer computer systems or peripherals based on widely accepted standards.)

IV. Suppliers Will Continue to Offer Differentiated High end Systems

Some observers suggest that the adoption of industry-standard 64-bit processors will result in a vast sea of undifferentiated high-performance workstations and servers, just as the use of an industry-standard 32-bit (x86) architecture led to the commoditization of desktop computers and low end servers. Insight 64 discounts the likelihood of such an outcome. Low end servers and desktop PCs differ from mid-range and high end servers in three key regards:

First, low end systems possess highly constrained functional requirements that are easily addressed using highly integrated solutions. Mid-range and high end multiprocessor systems address a more open-ended set of requirements. In their efforts to satisfy these requirements, designers make trade-offs that impact the scalability, expandability, cost and performance of the solutions they create. Even with these trade-offs, such solutions are costly to develop and expensive to produce.

Second, the low ASPs of low end servers discourage system suppliers from making substantial technology investments in these products. Conversely, the high ASPs of systems selling for more than $25,000 generate revenues that support the development of complex proprietary core logic chip sets. This is hardly a surprising result; the business models for mid-range and high end systems have long supported the development of proprietary core logic chip sets along with proprietary processors. In fact, as shown in Figure 3, these mid-range and high end systems account for more than 60 percent of all server revenues.

Third, low end servers sell in relatively high numbers, compared with the tiny unit sales of mid-range and high end systems. As shown in Figure 3, systems priced below $25,000 account for approximately 95 percent of all server units. This high unit volume allows third parties to develop and market core logic solutions on a profitable basis. The availability of these solutions allows system suppliers to outsource low end platform technology. Conversely, the relatively low unit volume of mid-range and high end systems makes it difficult for a third party to develop and profitably market the complex multi-chip core logic solutions.

All in all, third-party technology solutions have thrived at the low end of the server market, where the complexity of the development task is low and prospective unit sales of the solution are high. This phenomenon has commoditized the low end of the market. Third party technology solutions have failed to take hold at the high end of the market, where the complexity of the development task is high but the prospective unit sales of a solution are low, making it difficult for third parties to pursue a profitable business. This phenomenon assures that high end systems will continue to rely on proprietary core logic, and these core logic solutions will allow system suppliers to differentiate their systems in the market.

Figure 3
Distribution of Server Revenues and Shipments by System Price
(Four Quarters Ended March 2002)

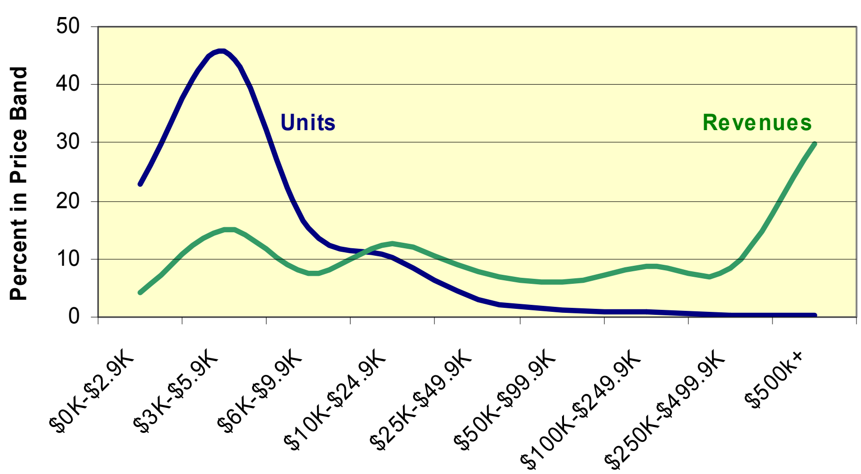

Source: IDC

V. Handicapping the Industry-standard Microprocessor Race

Throughout this document, Insight 64 has been explicit regarding its view that systems based on industry-standard 64-bit microprocessors will prevail over those based on proprietary processors. We have until now avoided the issue of *which* industry-standard 64-bit processor will prevail – Intel's Itanium Processor Family or AMD's forthcoming Opteron line. It is far too soon to make serious predictions in this regard, given that AMD has yet to launch its Opteron line, and little is known about the actual performance or the reception AMD may get from OEMs and end users. Nevertheless some conclusions can be drawn from information suppliers have already made public.

We begin with the observation that buyers of mid-range and high end systems are notoriously conservative with regard to their system selection criteria. They strongly prefer to purchase platforms they trust from suppliers they know. They skeptically evaluate new technology from known suppliers, and even more skeptically evaluate new technology from credible suppliers with whom they have not already established relationships. New technologies cannot climb this wall of doubt unless they offer dramatic economic advantages, or allow an organization to handle tasks that cannot be achieved with proven systems from established suppliers. This market segment's conservatism affects the reception accorded to both Itanium and Opteron, albeit in different ways. Insight 64 will apply the same criteria we used earlier in our assessment of proprietary processors (as a group) against industry-standard processors, (i.e., performance, scalability, system reliability, software availability and economics), to compare the outlook for the Itanium processor family and the Opteron line.

Performance: It is impossible to compare the performance of Intel and AMD offerings at this time. AMD hints that its initial Opteron processor will seize the lead with regard to processor-intensive benchmarks like SPEC's CPU2000, but the company has yet to release data that supports this view; indeed it has released no data whatsoever with regard to the chip's performance. Intel's Itanium 2 processor seized the SPEC CPU2000 performance crown at the time of its July 2002 launch, and clearly established its credentials as a viable contender in the 64-bit marketplace. When AMD's Opteron enters the market in 2003, it will likely face Intel's next-generation Itanium processor, code-named Madison, the first Itanium processor to be manufactured using Intel's 130nm process technology. Madison's 6MB L3 cache, combined with the higher clock frequencies enabled by a process shrink, will further boost its performance over the already fast Itanium 2. *Advantage: Too soon to tell.*

Scalability: Itanium and Opteron differ substantially with regard to scalability. Intel's Itanium processors use a high-bandwidth bus to communicate with the system's core logic. This approach allows OEMs to design core logic that matches the scalability requirements of their target markets, which may include as few as two and as many as 100 Itanium processors. The large (3MB) L3 cache on the Itanium 2 processor and the even larger (6MB) L3 cache on forthcoming Madison further enhance the Itanium processor's scalability. AMD's Opteron implementation differs substantially. Opteron processors incorporate core logic functions on the CPU die. AMD's approach simplifies the design of two-way and four-way configurations, but constrains OEM design options for larger configurations. Insight 64 believes end-users are unlikely to pursue large Opteron configurations in any event, given the conservative nature of buyers in this market and AMD's non-existent reputation as a supplier to this market. *Advantage: Itanium*

System Reliability: The reliability of large servers depends on a combination of hardware and software features. The designers of the Itanium processor had direct access to Hewlett-Packard's extensive experience with high end systems, and worked with the developers of HP's enterprise-class HP-UX Unix implementation to ensure the processor hardware included the facilities needed to implement a robust machine check architecture. AMD's Opteron design team lacked access to system developers actively engaged in high end reliability work. Nor did they have a complete software model, a la HP-UX, around which they could design their machine check architecture implementation. Many esoteric machine check features are required only in complex configurations with lots of processors. Most Opteron servers will be low end two or four way configurations, so the features AMD has included may prove more than adequate. *Advantage: Itanium*

Software Availability: Any comparison of the software available for Itanium and Opteron systems must take into consideration the dramatic difference in performance these systems exhibit when running 32-bit x86 software applications. Both Itanium and Opteron processors include the ability to execute 32-bit x86 programs. The designers of the Itanium processor did little to optimize this 32-bit mode of operation. HP estimates the Itanium 2 processor's x86 performance compares with that of a 350MHz Pentium Pro processor,[4] a competitive level of performance in 1998, but a bit slow by today's standards. Users may find Itanium's 32-bit performance acceptable for occasional use on undemanding applications, but the vast majority of applications need to be recompiled for 64-bit operation. Fortunately, Intel and its OEM partners have been working with ISVs for several years, and many key software packages have been ported, or will be soon ported to take advantage of Itanium's features. AMD claims 32-bit applications will run at least as fast on its Opteron processor as they do on AMD's fastest desktop processors, i.e., in excess of 2GHz. This may be a saving grace for AMD's 64-bit marketing efforts, since AMD has yet to demonstrate its ability to persuade ISVs to adapt their packages for the AMD 64-bit mode of operation.

The availability of some application software packages may depend on the availability of particular operating systems on the target hardware. Itanium has a clear advantage over Opteron in this regard. AMD has announced the planned availability of SuSe Linux, Red Hat's Advanced Server, and an unspecified version of Microsoft's 64-bit Windows operating system. Itanium OEMs and software developers have indicated their plans to port many versions of Linux, Microsoft's Windows .Net Server Edition, HP's enterprise-class HP-UX, OpenVMS and Nonstop Kernel to the Itanium platform. HP makes it especially easy for ISV's running on its HP-UX software on PA-RISC platforms to migrate their applications to Itanium platforms.

All in all, Insight 64 feels it must render a split decision when it comes to the issue of software availability. AMD has the advantage when it comes to the availability of 32-bit software running on its Opteron platform in 32-bit mode, but Intel has an advantage with regard to the availability of operating systems and software packages running in 64-bit mode on Itanium platforms. Advantage *(32-bit) Opteron; Advantage (64-bit): Itanium*

Economics: In our earlier assessment of the economic factors that impact proprietary and industry-standard processors, we identified two issues (R&D expenses and manufacturing costs) that relate to the processor supplier's cost of doing business, and two (switching costs and network effect benefits) that relate to end-user TCO expenses. From Insight 64's perspective, the supplier-side issues appear to be a wash, although AMD may gain a slight advantage because it can apply much of its 64-bit R&D to both its high volume desktop and lower volume server businesses. (Intel must maintain separate development efforts for its 32-bit desktop and server lines, and its 64-bit server line.) The advantage with regard to end-user TCO benefits will remain clouded until AMD demonstrates its ability to establish OEM partnerships with top-tier server suppliers, an area where Intel has already racked up more than two-score design wins for its 64-bit line and countless relationships for its 32-bit processors. AMD has yet to score even once with this elite group of OEMs. If AMD can establish such a relationship with a single top-tier supplier, then that supplier's customers will benefit with regard to their ability to mix-and-match 32-bit and 64-bit applications, but will be as constrained as users of proprietary processors in their ability to switch from one system supplier to another. If AMD succeeds in establishing relationships with multiple system suppliers, then the end-user customers of those suppliers can enjoy the TCO benefits that accrue to end-users able to switch system suppliers without affecting their operating system and application software environments. Insight 64 regards the minimization of switching costs as a key benefit of industry-standard processors. Thus, unless and until AMD demonstrates top-tier OEM support for its Opteron processor, we must grant the economic advantage to Intel. *Advantage: Itanium*

[4] HP, *Overview of the new Itanium® 2-based hp servers rx2600 and rx5670*, July 2002, p. 28

VI. Summary

This Insight 64 survey has reviewed millions of years of the geologic record, along with 30 years of microprocessor development. We conclude that before the end of the decade, industry-standard microprocessors will have pushed the proprietary predecessors that presently reign supreme into the fossil record. With regard to the much ballyhooed contest between 64-bit industry-standard processors from Intel and AMD, we conclude that both designs possess much potential. We envision the Itanium line as best equipped to tackle the high end of the market, controlled today by Sun, IBM and HP's RISC processors, while the Opteron line seeks a position at the high end of the workstation and low end server market. Regardless of the ultimate split between Itanium-based and Opteron-based systems, end-users will benefit from the high-performance workstation and server market's shift from proprietary to industry-standard systems.